



US005619497A

United States Patent [19]

Gallagher et al.

[11] **Patent Number:** 5,619,497[45] **Date of Patent:** Apr. 8, 1997[54] **METHOD AND APPARATUS FOR REORDERING FRAMES**[75] Inventors: **Brian Gallagher**, Marlboro; **Miklos A. Sandorfi**, Foxboro, both of Mass.[73] Assignee: **EMC Corporation**, Hopkinton, Mass.

[21] Appl. No.: 616,776

[22] Filed: **Mar. 15, 1996****Related U.S. Application Data**

[63] Continuation of Ser. No. 433,086, May 3, 1995, abandoned, which is a continuation-in-part of Ser. No. 363,392, Dec. 22, 1994.

[51] Int. Cl.⁶ **H04L 12/56**[52] U.S. Cl. **370/394**

[58] Field of Search 370/60, 94.1, 54, 370/60.1, 94.2, 94.3, 85.1, 112, 85.6, 61, 110.1, 85.13, 105.1, 82, 85.7, 74, 95.3, 95.1; 379/93, 94, 95; 340/825.5, 825.51, 825.52, 825.02; 395/825, 830, 250, 856, 860, 864, 871, 842, 850, 280, 427; 359/139

[56] **References Cited****U.S. PATENT DOCUMENTS**

Re. 32,327	1/1987	Biba et al.	370/124
4,451,916	5/1984	Casper et al.	370/16
4,654,874	2/1987	Fildes	379/93
4,763,317	8/1988	Lehman et al.	370/58
4,811,210	3/1989	McAulay	364/200
4,821,265	4/1989	Albal et al.	370/110
4,860,111	8/1989	Shimizu et al.	358/400

(List continued on next page.)

OTHER PUBLICATIONS

Seagate Product Specification, "ST15150FC Barracuda Disc Drive, Fibre Channel Arbitrated Loop Interface", Aug. 1995, cover sheet and pp. 1-75.

Seagate Technology, Inc., Twin Cities Operations, "Functional Description and Register Definition of the Fibre Channel Portion of Aurora", Aug. 1, 1994, pp. 1-59, marked Proprietary.

American National Standard for Information Systems, Fibre Channel, Physical and Signaling Interface (FC-PH), Rev. 4.3, Jun. 1, 1994, pp. i-xxiv, 1-388 and 11-19.

American National Standard for Information Systems, Fibre Channel Arbitrated Loop (FC-AL) rev. 4.4, Oct. 12, 1994, pp. i-x and 1-90.


International Business Machines Corporation, Serial Storage Chip Technical Reference, SSA-UIG/93-026 rev. 0, May 14, 1993, pp. i-viii and 1-245.

Primary Examiner—Douglas W. Olms**Assistant Examiner**—Dang Ton**Attorney, Agent, or Firm**—Weingarten, Schurgin, Gagnebin & Hayes LLP

[57]


ABSTRACT


A method and apparatus for coupling a plurality of channels of a communications network to a node. The node includes a plurality of ports, each adapted for connection to a corresponding channel and a system interface adapted for connection to a plurality of node clients in the form of host computers, peripheral devices, network interfaces, etc. The node includes a receiver and transmitter dedicated to each port and common circuitry for controlling and processing frames received and/or transmitted by the plurality of ports. A frame prioritization circuit forwards frames received by a selected port to a frame handler for processing and a frame routing circuit routes frames processed by the frame handler to at least one of the ports associated with at least one destination node. The node further includes a frame header buffer associated with each port for temporarily storing the header of each frame received by the respective port prior to forwarding of the header to the frame handler.

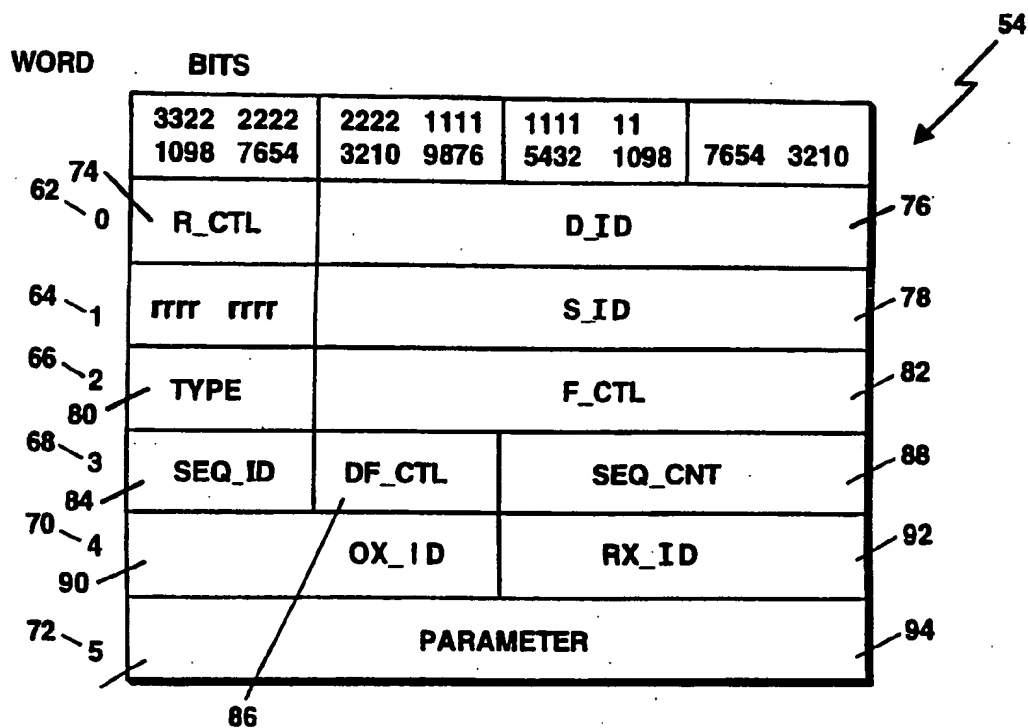
23 Claims, 16 Drawing Sheets**Best Available Copy**

U.S. PATENT DOCUMENTS

4,887,259	12/1989	Morita	370/60	5,237,571	8/1993	Cotton et al.	370/110.1
4,893,302	1/1990	Hemmady et al.	370/60	5,243,596	9/1993	Port et al.	370/94.1
4,899,333	2/1990	Roediger	370/60	5,245,339	9/1993	Cideciyan	341/95
4,942,574	7/1990	Zelle	370/85.15	5,255,264	10/1993	Cotton et al.	370/24
4,949,338	8/1990	Albal et al.	370/85.6	5,260,933	11/1993	Rouse	370/14
4,962,497	10/1990	Ferenc et al.	370/60.1	5,291,482	3/1994	McHarg et al.	370/60
5,027,349	6/1991	Thorne	370/85.1	5,295,137	3/1994	Jurkevich	370/54
5,123,109	6/1992	Hillis	395/800	5,303,078	4/1994	Brackett et al.	359/139
5,128,789	7/1992	Abramovitz	359/118	5,321,631	6/1994	Buhrke et al.	370/60
5,130,984	7/1992	Cisneros	370/94.1	5,329,524	7/1994	Paker et al.	370/58.1
5,142,672	8/1992	Johnson et al.	395/500	5,331,315	7/1994	Crosetto	340/825.02
5,150,357	9/1992	Hopner et al.	370/68.1	5,333,131	7/1994	Tanabe et al.	370/54
5,151,996	9/1992	Hillis	395/800	5,335,325	8/1994	Frank et al.	395/200
5,153,757	10/1992	Cloonan et al.	359/117	5,337,414	8/1994	Hashemi et al.	395/275
5,157,654	10/1992	Cisneros	370/60	5,345,446	9/1994	Hiller et al.	370/60.1
5,163,131	11/1992	Row et al.	395/200	5,351,236	9/1994	Pawelski	370/58.1
5,172,259	12/1992	Cloonan et al.	359/139	5,359,602	10/1994	Diaz et al.	370/85.8
5,177,739	1/1993	Basnuevo et al.	370/85.8	5,365,524	11/1994	Hiller et al.	370/94.2
5,195,091	3/1993	Farwell et al.	370/94.1	5,367,518	11/1994	Newman	370/54
5,197,064	3/1993	Chao	370/60	5,369,682	11/1994	Witsaman et al.	379/57
5,204,949	4/1993	Yasue et al.	395/200	5,371,852	12/1994	Attanasio et al.	395/200
5,222,085	6/1993	Newmann	370/60	5,378,943	1/1995	Dennard	326/68
5,226,125	7/1993	Balmer et al.	395/325	5,379,280	1/1995	Cotton et al.	370/62
5,233,607	8/1993	Barwig et al.	370/94.1	5,379,295	1/1995	Yonehara	370/94.1
5,235,690	8/1993	Beardsley et al.	395/425	5,381,405	1/1995	Daugherty et al.	370/54





**FIGURE 2B**

180

	SEQUENCE COUNT	DATA POINTER	DATA LENGTH
182a	<u>184a</u>	<u>186a</u>	<u>188a</u>
182b	<u>184b</u>	<u>186b</u>	<u>188b</u>
	⋮	⋮	⋮
182n	<u>184n</u>	<u>186n</u>	<u>188n</u>

FIGURE 5

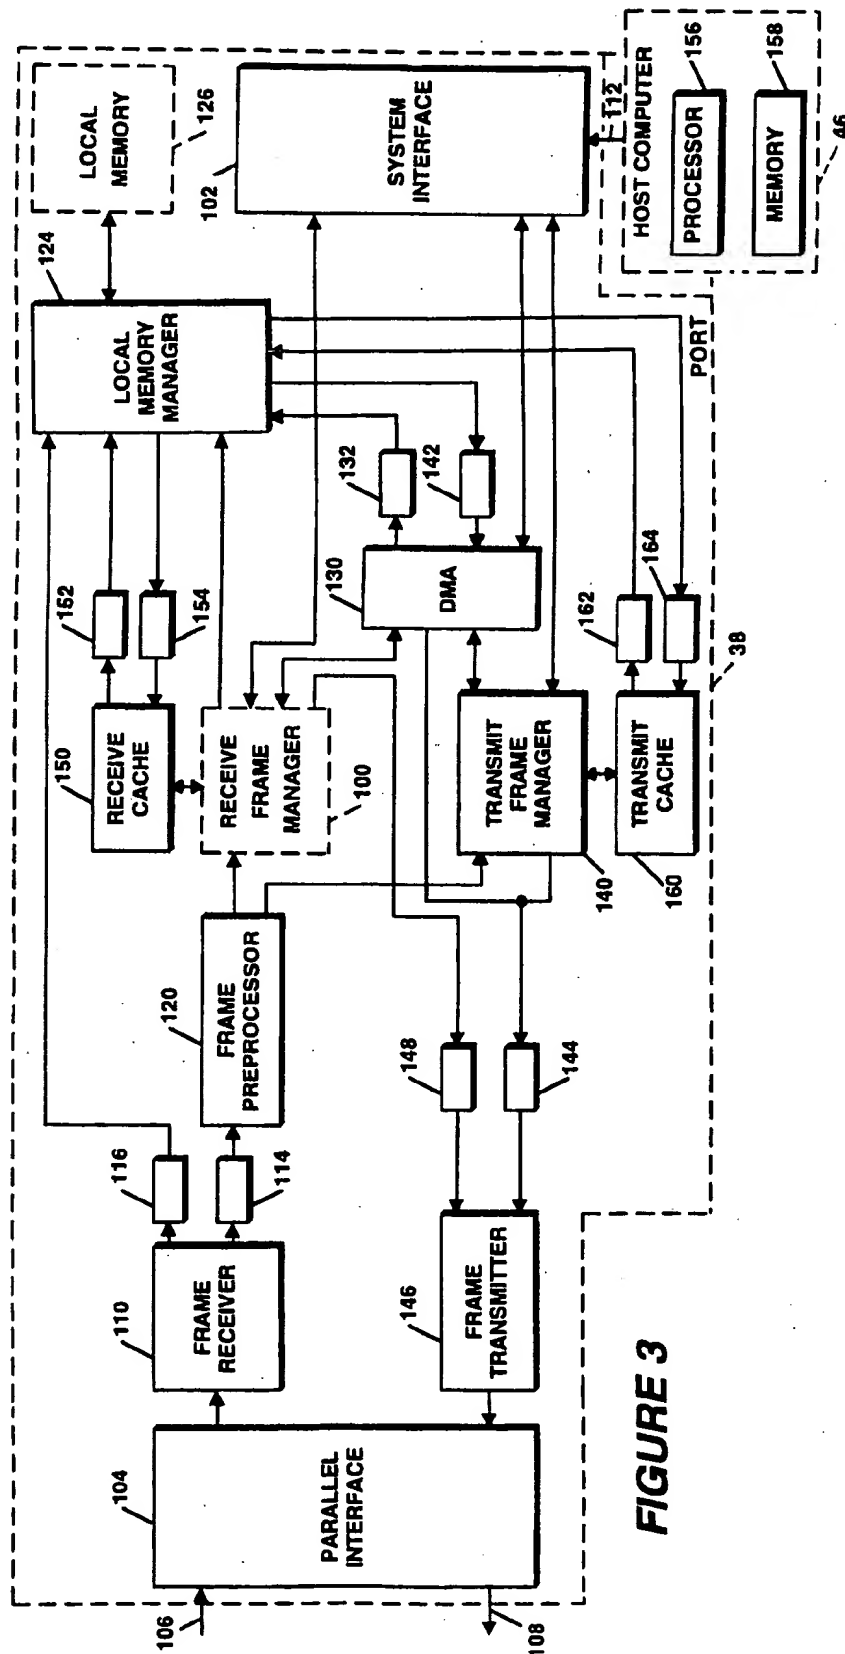


FIGURE 3




FIGURE 4


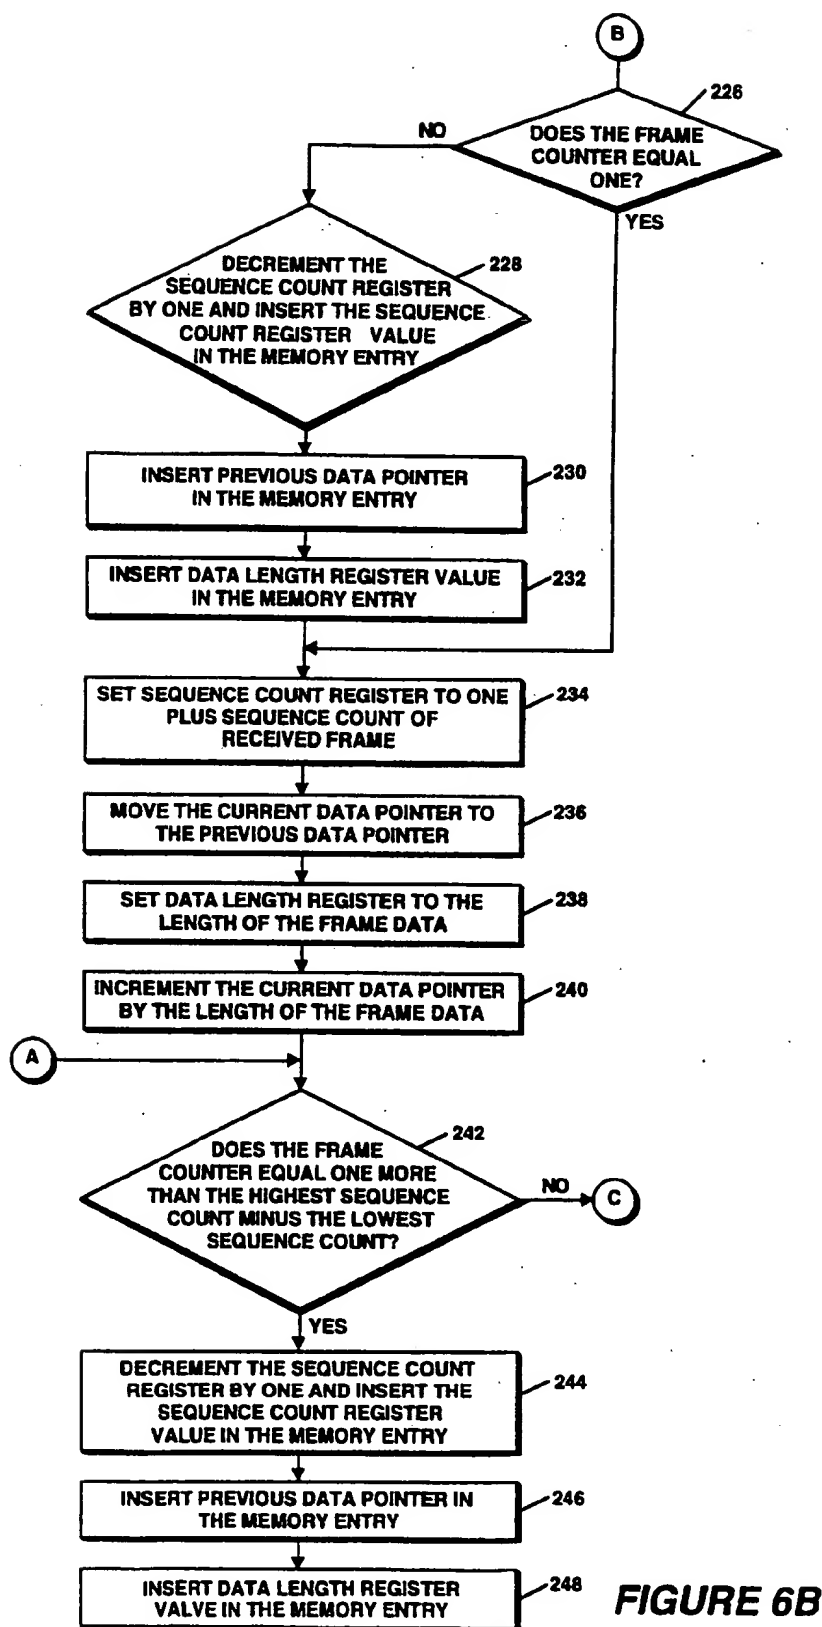


FIGURE 6A

**FIGURE 6B**




FIGURE 7A




FIGURE 7B





FIGURE 7C

250	0-127	DATA 56a OF FRAME 50a
	128-255	
	256-383	
	384-511	
252	512-639	DATA 56b OF FRAME 50b
	640-767	
254	768-895	DATA 56f OF FRAME 50f
	896-1023	
	1024-1151	
	1152-1279	
	1280-1407	
	1408-1535	
	1536-1663	
	1664-1791	
256	1792-1919	DATA 56c OF FRAME 50c
	1920-2047	
	2048-2175	
	2176-2303	
258	2304-2431	DATA 56d OF FRAME 50d
	2432-2559	
	2560-2687	
	2688-2815	
	2816-2943	
	2944-3071	
	3072-3199	
	3200-3327	
	3328-3455	
	3456-3583	
	3584-3711	
	3712-3839	
	3840-3967	
	3968-4095	
	4096-4223	
260	4224-4351	
262	4352-4479	DATA 56e OF FRAME 50e

FIGURE 8


180'



	SEQUENCE COUNT	DATA POINTER	DATA LENGTH
182a'	1	250	768
182b'	5	254	1024
182c'	4	256	2688

FIGURE 9

180"



	SEQUENCE COUNT	DATA POINTER	DATA LENGTH
182a''	1	270	256
182b''	0	272	512
182c''	5	274	1024
182d''	4	276	2688

FIGURE 11

270	→	0-127	DATA 56b OF FRAME 50b
		128-255	
272	→	256-383	DATA 56a OF FRAME 50a
		384-511	
		512-639	
		640-767	
274	→	768-895	DATA 56f OF FRAME 50f
		896-1023	
		1024-1151	
		1152-1279	
		1280-1407	
		1408-1535	
		1536-1663	
		1664-1791	
276	→	1792-1919	DATA 56c OF FRAME 50c
		1920-2047	
		2048-2175	
		2176-2303	
278	→	2304-2431	DATA 56d OF FRAME 50d
		2432-2559	
		2560-2687	
		2688-2815	
		2816-2943	
		2944-3071	
		3072-3199	
		3200-3327	
		3328-3455	
		3456-3583	
		3584-3711	
		3712-3839	
		3840-3967	
		3968-4095	
		4096-4223	
		4224-4351	
280	→	4352-4479	DATA 56e OF FRAME 50e
282	→		

FIGURE 10





FIGURE 12

**FIGURE 13**

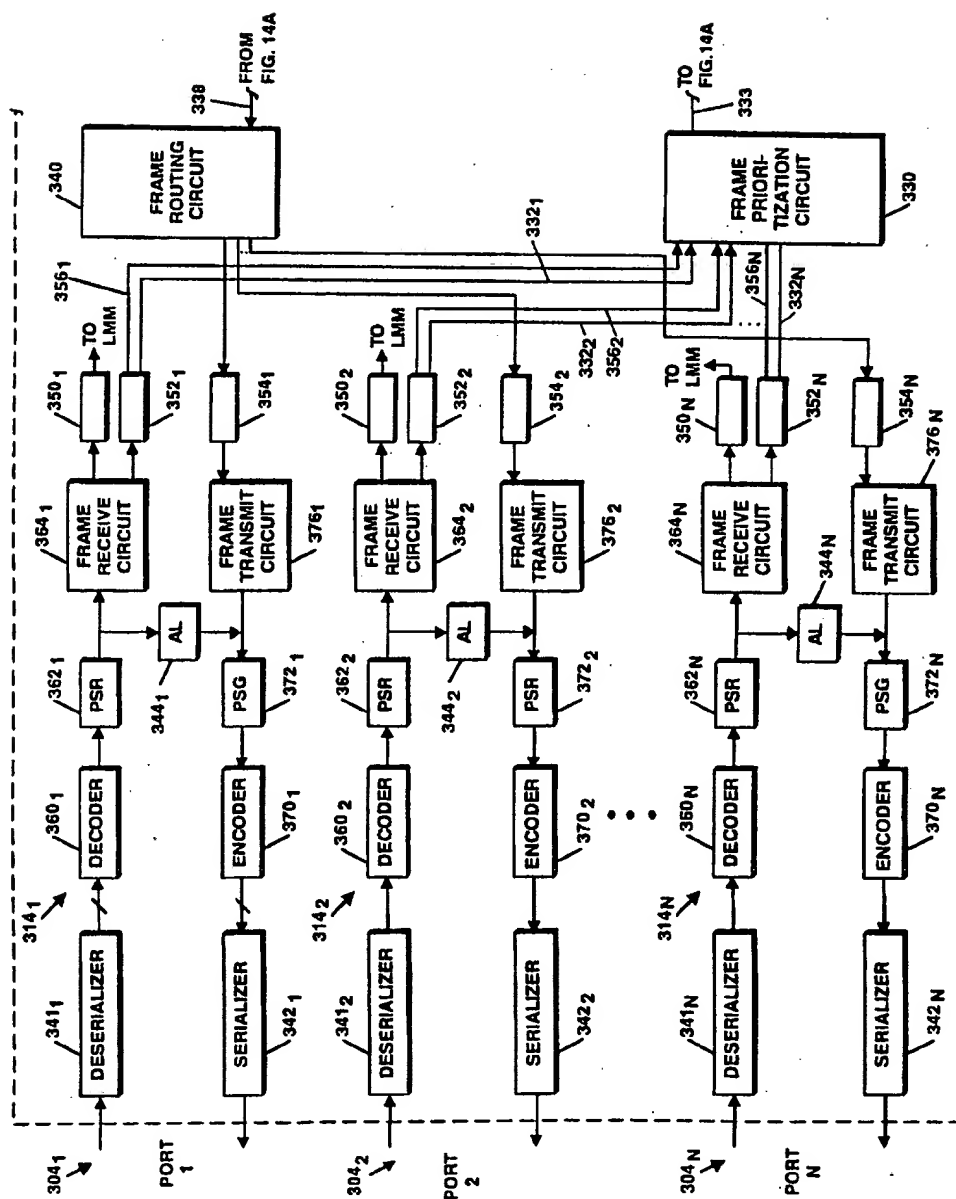


FIGURE 14

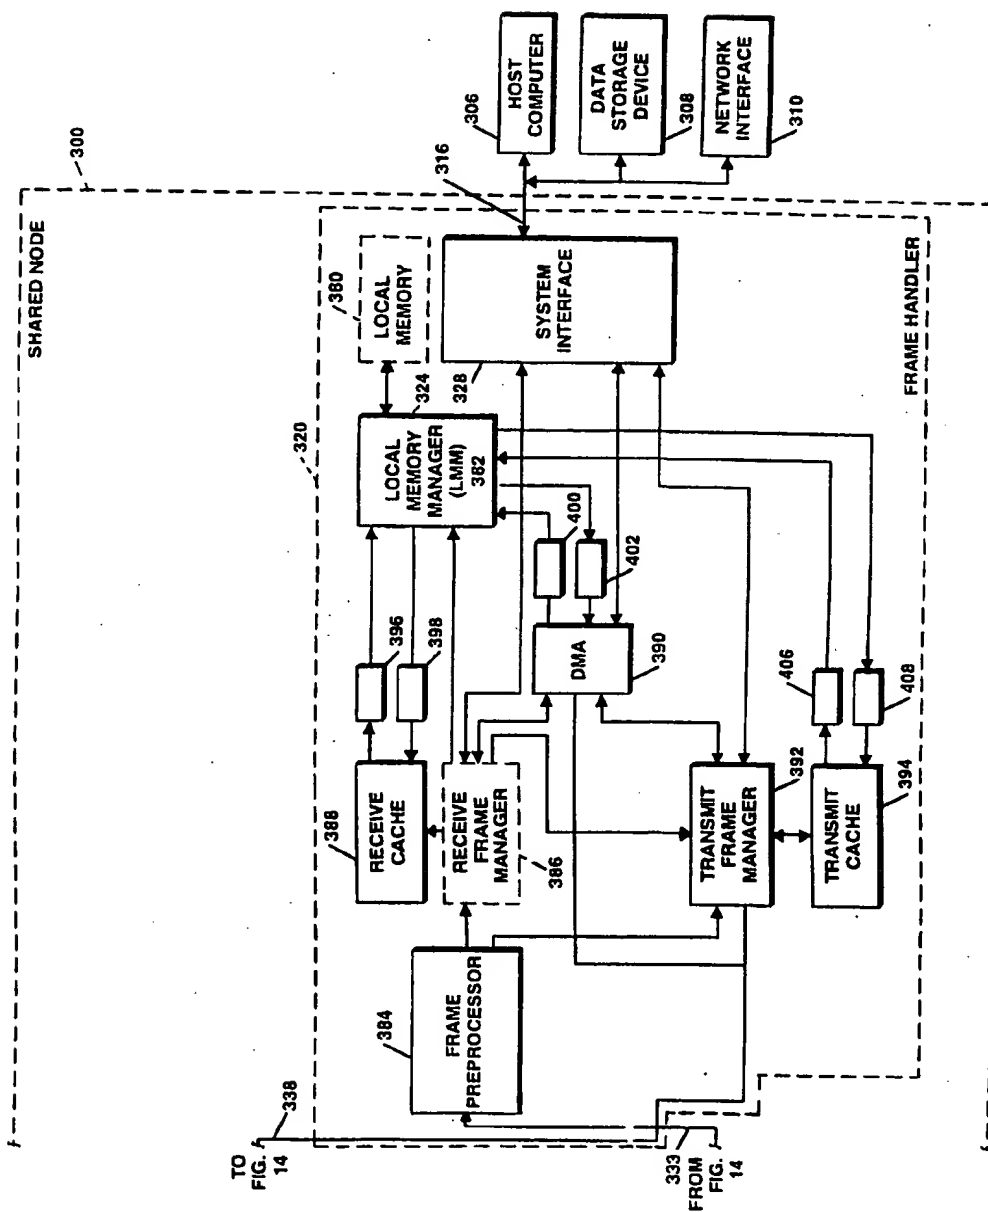




FIGURE 14A

**FIGURE 15****FIGURE 16**

METHOD AND APPARATUS FOR REORDERING FRAMES

This application is a continuation of application Ser. No. 08/433,086, filed May 3, 1995 entitled: DYNAMIC NODE SHARING APPARATUS AND TECHNIQUES which is a continuation-in-part application of U.S. Ser. No. 08/363,392, filed Dec. 22, 1994 entitled: METHOD AND APPARATUS FOR REORDERING FRAMES.

FIELD OF THE INVENTION

This invention relates generally to communications networks and more particularly, to apparatus and techniques for processing frames at a node having a plurality of ports.

BACKGROUND OF THE INVENTION

Switch networks, or fabrics, for transmitting packets between two or more nodes are known. Typically, a packet, sometimes referred to alternatively as a sequence, includes multiple frames, with each frame including a header and a payload, such as data. Switch fabrics include a plurality of switches arranged to permit transmission of frames over different paths, or channels. Each node has a port adapted for connection to a respective channel. One illustrative type of switch fabric is Fibre Channel, in which the nodes connected to the fabric are referred to as N_Ports.

Illustrative nodes are connected to various types of systems or node clients, such as work stations, data storage devices, or arrays, and network interfaces. A network interface node client is adapted for connection to a network, such as a local area network, thereby permitting communication between other nodes connected to the switch fabric and nodes connected to the network. Although generally a node is capable of both transmitting and receiving packets, in the transmission of a given packet, the node transmitting the packet is referred to hereinafter as the transmitting, or source node and the node receiving the packet is referred to hereinafter as the receiving, or destination node.

Each node includes circuitry for receiving frames from a source node over the respective channel, processing received frames, and transmitting processed frames to one or more destination nodes over the respective channel. Thus, each node has dedicated frame receiving, processing, and transmitting circuitry associated therewith.

Various techniques are employed for optimizing transmission of frames through a fabric in order to minimize congestion on a particular path and thus, to reduce transmission time. However, such transmission optimization may result in frames of a given packet arriving at the destination node in an order different from the order in which the frames are arranged in the packet prior to transmission.

One technique for reordering frames received at a node out of order is to generate and store a memory entry corresponding to each of the received frames, including a "sequence count" identifying the sequential position of the frame within its associated packet. Once all of the frames of a packet are received, a processor executes a program to sort the memory entries by sequence count. However, since the memory includes as many entries as there are frames in the packet, this technique tends to be inefficient in terms of both memory requirement for storing the entries and processing time for sorting the entries.

SUMMARY OF THE INVENTION

In accordance with the invention, a receive node is provided with a port for reordering frames of a packet

received at the port out of order in a manner that requires less memory space and processing time than heretofore required. Note that in the case where a node has only one port associated therewith, the terms "node" and "port" are used interchangeably. The port includes a frame manager which receives frames and generates a reassembly table in an associated memory including only entries corresponding to frames received out of sequential order with respect to a previously received frame (and a last entry indicating that the last frame of the respective packet has been received). In one embodiment, each such memory entry includes a sequence count, identifying the previously received frame, a pointer to a location in memory at which data associated with a set of one or more previously, consecutively received frames are stored, and a data length value representing the length of data associated with the set of frames.

The reassembly table entries are processed to reorder the frames of the received packet. More particularly, the port is in communication with a system which includes a processor for processing entries of the reassembly table in ascending sequence count order to reorder the packet frames. Illustrative systems include a host computer, a disk array, and a network.

In one embodiment, the receiving node processes frames of multiple packets simultaneously and the frame manager generates multiple reassembly tables, each one corresponding to a processed packet. To this end, a context is stored for each processed packet and includes fields used to generate memory entries for non-consecutively received frames of the respective packet.

A method for reordering a plurality of frames of a packet received by a node in non-consecutive order includes the steps of receiving a frame, comparing the sequence count of the received frame to a stored sequence count, and providing an entry in memory if the sequence count of the received frame is not equal to the stored sequence count. The stored sequence count corresponds to the sequence count of a previously received frame incremented by one, or to an initialized value when the received frame is the first frame to be received. The method may further include the steps of maintaining a value representative of a length of data associated with a set of one or more prior consecutively received frames and providing the data length value in the memory entry. A pointer may also be provided in the memory entry to a location in memory at which data associated with the set of prior consecutively received frames are stored. The number of frames received for a given packet are counted and the reordering method is terminated when the number of counted frames is equal to the number of frames in the given packet.

With this arrangement, a port of a receiving node generates a reassembly table which permits packet frames to be reordered without requiring significant memory space and without imposing significant processing delays heretofore associated with frame reordering. These efficiencies are achieved by generating a reassembly table containing only entries indicating non-consecutive breaks of received frames, as opposed to storing entries for each received frame. Since the reassembly table contains fewer entries than received frames, and fewer entries than conventional reordering techniques, the memory requirement is reduced and the time associated with processing the memory entries to reassemble the packet is reduced.

In accordance with a further embodiment of the invention, a shared node is provided with multiple ports adapted for connection to channels of a switch fabric. In the case where

3

the shared node has a plurality of ports associated therewith, the term "node" refers to the entire switch fabric interconnection circuit whereas, the term "port" refers to the portion of the node used to connect the node to a particular channel of the switch fabric. Each port of the shared node is capable of both receiving and transmitting frames from and to the channel respectively and has a frame header buffer associated therewith for temporarily storing the header of frames received by the port. A plurality of node client systems are connected to the shared node via a bus, with each such node client logically associated with a particular channel. Illustrative node client systems include host computers, data storage devices and network interfaces. The shared node includes a frame prioritization circuit, or multiplexer, for selecting frames received by the ports for forwarding to a frame handler for processing. Also provided is a frame routing circuit, or demultiplexer, for routing frames processed by the frame handler to at least one of the ports for transmission to one or more destination nodes.

With this arrangement, a common shared portion of node circuitry is provided for receiving, processing, and transmitting frames associated with a plurality of network channels. In this way, the cost of providing a node is reduced, since frame processing circuit duplication for each channel is avoided. Moreover, utilization of shared resources, such as local memory, by multiple ports, serves to efficiently distribute the cost of such resources among the ports. Also, the circuit board area occupied by each node is reduced by sharing circuitry among multiple ports.

The frame prioritization circuit selects received frames for forwarding to the frame handler in accordance with one of several prioritization schemes. In one embodiment, the frame prioritization circuit selects frames in accordance with a priority assigned to each port at system initialization. The preassigned port priority may be based on various criteria, such as expected frame traffic over a particular channel or the type of node client associated with a particular channel. In another embodiment, the frame prioritization circuit selects frames for forwarding to the frame handler based on a predetermined order of the ports, such as a "round robin" sequence of the ports or based on the order in which the ports receive frames. Alternatively, and in accordance with an adaptive scheme for selecting frames for forwarding to the frame handler, the frame prioritization circuit determines how many frame headers are stored in the frame header buffer associated with each port and forwards frames based on which port has the most frame headers stored in the respective frame header buffer.

The frame prioritization scheme can be tailored to permit manipulation of the operating parameters specified by each node to every other node at system initialization. More particularly, although the frame header buffer associated with each port has a fixed size, the "credit" (i.e. frame receiving capacity specified to other nodes) can indicate that greater than actual buffer capacity is available for receiving frames when the frame prioritization scheme is tailored to ensure that the frames received at that particular port are rapidly forwarded to the frame handler. With this arrangement, a port can be permitted to accommodate a higher frame traffic level than otherwise permissible by tailoring the frame prioritization scheme to provide the port with a higher rate of forwarding frames to the frame handler. Thus, use of the shared node permits frame traffic allocation among the multiple channels serviced by the node, thereby permitting optimization of node performance.

The frame routing circuit includes a control circuit for determining a destination node specified in a processed

4

frame and a look-up table. The table includes a list of known nodes connected to the switch fabric and an associated port of the shared node in communication with each listed node. In operation, the control circuit accesses the table to locate the destination node specified in the processed frame, reads the identifier of the port associated with the destination node, and routes the frame to the corresponding port for transmission to the destination node.

Also described is a method for processing frames received by a receive port, including the steps of receiving frames at a plurality of ports from a corresponding plurality of channels, with the ports being capable of receiving frames substantially simultaneously with respect to one another, selecting a frame received by one of the ports for processing by a frame handler, forwarding the selected frame to the frame handler and routing the processed frame to a selected port associated with a destination node. In one embodiment, the frame selecting step includes determining a priority assigned to the ports and selecting frames in accordance with a descending order of port priority. In an adaptive frame selection embodiment, the frame selecting step includes selecting frames in accordance with the number of frames received by each port. The frame routing step includes determining the identity of a destination node specified in a processed frame, accessing a look-up table to determine a port associated with the determined destination node, and routing the processed frame to the port associated with the destination node.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 is a diagram of an illustrative switch fabric interconnecting a plurality of nodes having ports in accordance with the invention;

FIG. 2A shows an illustrative frame format for transmission over the switch fabric of FIG. 1;

FIG. 2B shows an illustrative format of a header associated with the frames of FIG. 2A;

FIG. 3 is a block diagram of an illustrative port of FIG. 1;

FIG. 4 is a block diagram of the Receive Frame Manager of FIG. 3;

FIG. 5 shows an illustrative format of a reassembly table in accordance with the invention;

FIGS. 6A and 6B show a flow chart of an illustrative process by which the port of FIG. 3 reorders frames;

FIG. 7A shows an illustrative packet transmitted by a transmitting node of FIG. 1;

FIG. 7B shows one illustrative order in which the frames of the packet of FIG. 7A are received by a receiving node;

FIG. 7C shows another illustrative order in which the frames of the packet of FIG. 7A are received by a receiving node;

FIG. 8 shows an illustrative portion of system memory allocated to store the received frames of FIG. 7B;

FIG. 9 shows a reassembly table containing entries associated with the received frames of FIG. 7B;

FIG. 10 shows an illustrative portion of system memory allocated to store the received frames of FIG. 7C; and

FIG. 11 shows a reassembly table containing entries associated with the received frames of FIG. 7C;

FIG. 12 is a diagram of a switch fabric interconnecting a plurality of nodes in accordance with a further embodiment of the invention;

FIG. 13 is a block diagram of the shared node of FIG. 12;

FIGS. 14 and 14A are a detailed block diagram of the shared node of FIG. 12;

FIG. 15 is a diagram of the frame prioritization circuit of the shared node of FIGS. 12-14A; and

FIG. 16 is a diagram of the frame routing circuit of the shared node of FIGS. 12-14A.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to FIG. 1, a switch fabric 10, including a plurality of switches 12, 14, 16, 18, and 20, interconnects a plurality of nodes 22, 24, 26. The switches 12-20 are arranged to permit transmission of packets, frame by frame, between the nodes 22-26 over a plurality of paths 28, 30, and 32. More particularly, the switches 12-20, sometimes referred to as routers, selectively direct frames over an optimized path 28-32, such as the path having the least traffic thereon, or the fastest communication link by various known techniques. While a switch fabric 10 is illustrated, the network 10 may comprise any point-to-point interconnection network.

Each of the nodes 22, 24, 26 includes a respective port 36, 38, 40 (referred to sometimes as an N_Port) connected to a switch 12-20 of the fabric 10 for receiving frames from the respective switch and for transmitting frames to the respective switch, as will be discussed. Note that where a node has only a single port associated therewith, as is the case for nodes 22, 24, 26 and respective ports 36, 38, 40, the terms "node" and "port" are used interchangeably. Various types of nodes may be connected to the fabric 10; such as, a data storage node 22, including a data storage device 44 in communication with port 36, a workstation node 24, including a host computer 46 in communication with port 38, and a network node 26 including a network interface 48 in communication with port 40. Such a network interface 48 is further adapted for connection to a network (not shown) having additional nodes connected thereto. It will be appreciated that other types of nodes may be connected to the fabric 10 and that additional nodes (not shown) may be connected to the fabric 10.

The node ports 36-40 include apparatus for reordering frames of a packet received by the respective node in an order other than the order in which the frames are arranged in the packet as transmitted, as will be discussed. In one embodiment, the transmitted packets have a format defined by a Fibre Channel standard (in a specification entitled "FIBRE CHANNEL PHYSICAL AND SIGNALING INTERFACE STANDARD" (FC-PH) Rev. 4.3, X3, 230-1994) and shown in FIGS. 2A and 2B for transmission over a Fibre Channel fabric. However, it will be appreciated that the invention may be used with various types of fabrics 10 which permit frames of various formats to be received at a destination node in an order different than the order in which the frames were transmitted.

Referring also to FIG. 2A, an illustrative Fibre Channel packet 34 is shown to include a plurality of frames 50a-n spaced by idle words. Considering the format of exemplary frame 50a, each frame includes a Start-of-Frame (SOF) field 52a having a length of four bytes and indicating the class of service requested for the transmission and whether or not the frame 50a is the first or last in the packet 34 as transmitted.

More particularly, a Fibre Channel fabric supports three classes of packet transmission service. In one such class of service, a packet is transmitted over a single path 28-32, thereby precluding frames of the packet from arriving at a receiving node out of order. The remaining two classes of packet transmission service permit individual frames of a packet to be transmitted over different paths 28-32, and differ from one another in that, one such class of service requires that an acknowledgement, indicating that a frame was properly received, be sent by the receiving node to the transmitting node; whereas the other such class of service does not require an acknowledgement. Since frames of a packet can be transmitted over different paths 28-32 in the latter two classes of Fibre Channel service, such frames can arrive at the receiving node in an order different than the order in which the frames were transmitted. That is, while the transmitted packet includes a plurality of sequentially ordered frames (as specified by a sequence count contained within each frame), frames may be received at the receiving node out of this sequential order.

A frame header 54a, twenty-four bytes in length, follows the SOF field 52a, and will be described further below in conjunction with FIG. 2B. A data field 56a having a length of between 0 and 2112 bytes follows the frame header 54a, and, in turn, is followed by a four byte Cyclic Redundancy Check (CRC) field 58a which is used to verify the integrity of the frame 50a. A four byte End-of-Frame (EOF) field 60 follows the CRC field 58a and indicates whether or not the frame 50a is valid or corrupted.

Referring also to FIG. 2B, the header 54a of exemplary frame 50a is shown to include six words 62, 64, 66, 68, 70, and 72, each thirty-two bits long. The first byte of the first word 62 provides an R_CTL field 74 which defines the frame type as containing either a data field length of zero (i.e., a link control frame) or a data field length of between zero and 2112 bytes (i.e., a data frame). The remaining three bytes of the first word 62 provide a D_ID field 76, identifying the destination node and three bytes of the second word 64 provide an S_ID field 78, identifying the source node. The first byte of the third word 66 provides a TYPE field 80 which, in conjunction with the R_CTL field 74, specifies the protocol of the data associated with the frame. The remaining three bytes of the third word 66 provide an F_CTL field 82 which contains control information regarding the frame, including a one bit flag, at bit nineteen, indicating whether the frame 50a is the last frame of the packet 34 according to the order in which the frames were transmitted. The first byte of the fourth word 68 provides a SEQ_ID field 84 which contains an identifier of the sequence, or packet 34 with which the frame 50a is associated and is followed by a DF_CTL field 86 which specifies the presence of optional header bits at the beginning of the data field 56a. A sixteen bit SEQ_CNT field 88 is also provided in the fourth word 68 and contains a unique identifier of the frame 50a within the packet 34, referred to hereinafter as the sequence count. Note that where the SOF field 52a of a received frame 50 indicates that the frame is the first of the associated packet based on the order in which the frames appear in the packet as transmitted, then the sequence count of the frame provides a lowest sequence count, for use in a manner described below. Similarly, where bit nineteen of the F_CTL field 82 of a received frame indicates that the frame is the last in the respective packet as transmitted, then the sequence count of that frame provides a highest sequence count, for use in a manner described below.

The fifth word 70 of the header 54 includes an OX_ID field 90 and an RX_ID field 92. The OX_ID and RX_ID

fields 90, 92 carry an exchange identifier provided by the source node and the destination node, respectively. The sixth word 72 of the frame header 54a contains a parameter field 94, the content of which varies in accordance with whether the frame 50a is a link control frame or a data frame.

Referring to FIG. 3, the ports 36, 38, 40 will be described in conjunction with illustrative port 38 which, along with host computer 46, comprises workstation node 24 (FIG. 1). It will be appreciated however that the ports 36-40 associated with each of the nodes 22-26 have like construction and are adapted for connection to a system, such as a host computer or data storage device, having a processor 156 and a system memory 158.

Port 38 includes a Receive Frame Manager 100 which is responsive to frames 50a-n received by the port 38 for providing an entry in a memory 126 for each such frame 50a-n received out of the sequential order in which the frame 50a-n appears in the transmitted packet 34, as will be described. A list of each such memory entry associated with a given packet 34 provides a reassembly table for use in reordering the frames 50a-n to reassemble the packet 34. More particularly, the system processor 156 accesses the reassembly table, processing the entries contained therein, to reorder the received frames 50a-n.

The port 38 includes a parallel interface 104 connected to a serializer/deserializer (not shown) which is part of the node 24, positioned between the switch 18 and the port 38. The serializer/deserializer converts the serial Fibre Channel signals into digital signals for processing by the port 38. Thus, the parallel interface 104 receives parallel signals via a receipt signal line 106 and transmits parallel signals via a transmit signal line 108. The parallel interface 104 clocks incoming words of frames 50a-n from the receipt line 106 into a Frame receiver 110. As a frame 50a-n (referred to hereinafter generally as frame 50, including an SOF field 52, a header 54, data 56, a CRC field 58 and an EOF field 60) is received, the Frame receiver 110 separates the frame header 54 from the data 56 for further processing. The frame data 56 is transferred to a Frame Data FIFO 116 and the remainder of the frame 50, including the frame header 54, is transferred to a Frame Header FIFO 114. The frame header 54 is processed by a Frame preprocessor 120 connected to the Frame Header FIFO 114; whereas, the frame data 56 is transferred, via a Local Memory Manager 124 to local memory 126 for temporary storage. Note that while the local memory 126 is shown to be contained within the port 38, memory 126 may, alternatively, be located external to the port 38.

The Frame preprocessor 120 checks the CRC field 58 of the received frame 50 to determine whether the frame has been accurately received. In the event that the frame 50 has not been accurately received, the frame 50 is discarded. The Frame preprocessor 120 additionally determines whether the D_ID field 76 of the frame header 54 corresponds to the node 24 associated with the receiving port 38. In the event that the frame destination node specified in header field 76 does not match that of the receiving node 24, a reject frame is routed back to the fabric, for the first two classes of service, or discarded for the third class of service. After processing the frame header 54, the Frame preprocessor 120 transfers the frame header 54 to the Receive Frame Manager 100.

The Receive Frame Manager 100 controls the transfer of received frame data 56 from the local memory 126 to the system memory 158. More particularly, the Receive Frame Manager 100 programs a DMA controller 130 to move data

56 out of the local memory 126 and transfer such data to the system memory 158. The DMA controller 130 achieves this transfer of data 56 via a Write Data FIFO 132, a Read Data FIFO 142, the Local Memory Manager 124, and a System Interface 102. System Interface 102 permits communication between the host computer 46 and the port 38 via a system bus 112, by translating signals internal to the port 38 into signals compatible with the host computer 46.

When the port 38 transmits a frame 50 to the respective switch 18 (FIG. 1), a Transmit Frame Manager 140 programs the DMA controller 130 to read data 56 out of the system memory 158. The data 56 read by the DMA controller 130 is enqueued in a Frame Data FIFO 144 for transmission by a Frame transmitter 146. The Frame transmitter 146 transmits the header 54 followed by the frame data 56. The header 54 may be provided by the Transmit Frame Manager 140 or, alternatively, may be provided directly by the host computer 46. The Frame transmitter 146 is additionally responsive to a Response FIFO 148 which contains entries corresponding to acknowledgements provided by the Receive Frame Manager 100 that a frame was accurately received, as required by the first and second classes of a Fibre Channel service discussed above.

A Receive cache 150 is provided in communication with the Receive Frame Manager 100 and the Local Memory Manager 124 via a local memory Write Data FIFO 152 and a local memory Read Data FIFO 154. The Receive cache 150 stores a context associated with each packet processed by the port 38 for use by the Receive Frame Manager 100 in reordering packet frames, as will be described. Similarly, a transmit cache 160 is provided in communication with the Transmit Frame Manager 140 and the Local Memory Manager 124 via a local memory Write Data FIFO 162 and a local memory Read Data FIFO 164, as shown. The transmit cache 160 functions in a manner similar to the receive cache operation described below, albeit in the transmission of frames as opposed to the receipt of frames.

Referring also to FIG. 4, the Receive Frame Manager 100 includes a Receive Sequence state machine 170 and a Control state machine 174, both of which are in communication with the Frame preprocessor 120 (FIG. 3). The Receive Sequence state machine 170 is further in communication with the System Interface 102, the Receive cache 150 and the Local Memory Manager 124, as shown. The Control state machine 174 is connected to the DMA controller 130 and to the Response FIFO 148.

The Receive Frame Manager 100 additionally includes a current context 178, the contents of which are particular to the packet 34 associated with a currently processed frame 50. More particularly, each port 38 is capable of processing multiple packets 34 at any given time. When a frame 50 is received by the port 38, the packet 34 in which it resides is determined and the contents of the current context 178 are loaded with contents associated with that packet 34, as will be described.

The current context 178 includes registers 118 for maintaining information about the respective packet 34. Specifically, the context registers 118 include a Sequence Count register 122 which, in the illustrative embodiment, is sixteen bits long and stores a value equal to the sequence count of a next expected frame. To this end, when a frame 50 is received by the port 38, the Sequence Count register 122 is modified to contain the sequence count of the received frame 50 plus one. A Data Length register 128 is also provided in the context 178 and contains a value corresponding to the length of data associated with a set of one or more prior

consecutively received frames. In the illustrative embodiment, the Data Length register 128 is thirty-two bits long. The context registers 118 further include a Frame counter 134 for maintaining a sixteen bit count of the total number of frames 50a-n of a given packet 34 received by the port 38. A forty-four bit Previous Data Pointer register 136 is provided for maintaining a pointer, referred to hereinafter as the PREVIOUS DATA pointer to a location in the system memory 158 beginning at which is stored. Also provided in the context 118 is a CURRENT DATA Pointer register 138 which maintains a pointer, referred to hereinafter as the CURRENT DATA pointer, to a next available location in the system memory 158 for storing incoming frame data 56.

Upon receipt of a frame 50, the Control state machine 174 determines whether the type of service associated with a received frame 50 requires an acknowledgement to be sent to the source node indicating whether the frame was properly received, as specified in the SOF field 52 of the frame 50. If the class of transmission service requires an acknowledgement, the Control state machine 174 generates the acknowledgement and enques such acknowledgement in the Response FIFO 148.

As noted above, the Receive Frame Manager 100 programs the DMA controller 130 to read frame data 56 out of the local memory 126 and transfer such data to the system memory 158 via the system bus 112. Specifically, the Control state machine 174 programs the DMA controller 130 to read the data 56 into locations of the system memory 158 starting at the location pointed to by the CURRENT DATA pointer maintained in a CURRENT DATA Pointer register 138.

Once a frame 50 is received and the header 54 transferred to the Receive Frame Manager 100, the Receive Sequence state machine 170 generates a tag associated with the received frame 50 for comparison to contexts stored in the Receive cache 50. More particularly, the tag includes the source identification from the S_ID field 78 of the frame header 54, the sequence identification from the SEQ_ID field 84 of the header 54, and the OX_ID field 90 from the header 54. If the generated tag matches a tag stored in the Receive cache 150, then the contents of the matching context maintained in the cache 150 are loaded into the current context 178 maintained in the Receive Frame Manager 100 and the packet 34 containing the received frame 50 is referred to as a solicited transmission. Alternatively, if the generated tag does not match a context stored in the Receive cache 150, then the Receive Sequence state machine 170 prompts the system processor 156 for initialized contents for the current context 178 and the packet 34 associated with the received frame 50 is referred to as an unsolicited transmission.

Referring to FIG. 5, an illustrative format of a reassembly table 180 maintained in the Local Memory 126 for a packet 34 is shown. Recall that a reassembly table 180 is stored for each processed packet, such as packet 34. The table 180 includes one entry 182a-182n-1 for each frame of the packet 34 received non-consecutively with respect to a previously received frame and a last entry 182n corresponding to the last frame of the packet to be received.

Each table entry 182a-n includes a respective sequence count field 184a-n which is sixteen bits long, a respective data pointer field 186a-n which is forty-four bits long, and a respective data length field 188a-n which is thirty-two bits long. The sequence count field 184a-n contains the sequence count of the last consecutively received frame and is provided by the Sequence Count register 122 of the

current context 178. The data pointer field 186a-n contains a pointer to a location in the system memory 158 beginning at which sequential frame data is stored until a received discontinuity and is provided by the Previous Data Pointer register 136 of the current context 178. That is, the DATA POINTER points to the first location at which a set of one or more consecutive frames, received prior to the process, non-consecutively received frame, is stored. The data length field 188a-n contains a value indicating a length of data received by the port 38 since receipt of a last non-consecutively received frame (i.e., the length of data associated with the set of prior, consecutively received frames) and is provided by the Data Length register 128 of the current context 178.

Referring to FIGS. 6A and 6B, flow diagrams show an illustrative process by which frames 50a-n received out of sequential order by the port 38 are reordered. In step 200, the port 38 is powered up or connected to the fabric 10. The port 38 is initialized in steps 202-208 by initializing the contents of the current context registers 118 maintained in the Receive Frame Manager 100. Specifically, the Data Length register 128 is set to zero in step 202, the Sequence Count register 122 is set to the expected first sequence count in step 204, the PREVIOUS DATA pointer in the Data Pointer register 136 is set to the value of the CURRENT DATA pointer in the CURRENT DATA Pointer register 138 in step 206 and the Frame counter 134 is set to zero in step 208. Note that where a particular transmission includes a single packet, the Sequence Count register 122 is initialized to a value of zero. However, where a streamed sequence, including more than one packet is transmitted, the Sequence Count register 122 will be initialized to a value indicating the next sequential sequence count after the highest sequence count of a previously received packet of the streamed sequence.

Having initialized the current context 178, it is determined, in step 210, whether a frame 50a-n has been received by the port 38. Step 210 is repeated until a frame 50a-n is received. When a frame 50a-n, such as illustrative frame 50a, is received, the context 178 of the Receive Frame Manager 100 is loaded (unless the received frame 50a is the first frame of a packet in which case the registers remain in their initialized states provided in steps 202-208). More particularly, in step 212, the frame tag described above is generated by the Receive Sequence state machine 170 and is compared to tags stored in the Receive cache 150. If the generated tag is found to match a stored tag, then the contents of the current context registers 118 are loaded from the Receive cache 150 in accordance with registers associated with the matching tag. Alternatively, if the generated tag does not match a tag stored in the Receive cache 150 (i.e., if the received frame 50a is the first frame associated with a packet to be received), then the system processor 156 initializes the current context registers 118, in the same manner as the context 178 is initialized after system power up in steps 202-208.

Thereafter, in step 214, the frame data such as data 56a of received frame 50a, is stored in the system memory 158 at the location pointed to by the CURRENT DATA pointer (i.e., after having been temporarily stored in local memory 126). In step 216, the Frame counter 134 is incremented by one. In step 218, it is determined whether the sequence count of the received frame 50a is equal to the sequence count stored in the Sequence Count register 122. If the two sequence counts are equal, then the received frame 50a is consecutive with respect to the previously received frame (or, is the first frame of a packet to be received and is also the first sequential frame of the packet as transmitted) and no

entry is made in the reassembly table 180. Alternatively, if it is determined in step 218 that the sequence count of the received frame 50a is not equal to the sequence count stored in the Sequence Count register 122, then the received frame 50a is non-consecutive with respect to a previously received frame (or is the first frame of a packet to be received but is not the first sequential frame of the packet as transmitted) and an entry is made in the reassembly table 180.

Consider first the case where the sequence count of the received frame 50a is equal to the value in the Sequence Count register 122. In this case, no entry is made into memory and the current context registers 118 are reset in preparation for receipt of a subsequent frame in steps 220, 222, and 224. Specifically, the value in the Sequence Count register 122 is incremented by one in step 220, in preparation for comparison to the sequence count of a subsequently received frame. The value in the Data Length register 128 is incremented by the length of the received frame data 56a in step 222. In step 224, the CURRENT DATA pointer in register 138 is incremented by the length of the received frame data 56a, to indicate that the next available memory location is after the stored data 56a.

In the case where the sequence count of the received frame 50a is not equal to the value in the Sequence Count register 122, then step 226 is performed (FIG. 6B) after step 218, in which it is determined whether the Frame counter 134 is at a value of one. If it is determined that the Frame counter 134 is at one, thereby indicating that the received frame 50a is the first frame of the packet 34 received, then step 234 is next performed and no entry is made into the reassembly table 180. If, on the other hand, it is determined in step 226 that the value of the Frame counter 134 is not equal to one, then the non-consecutively received frame is not the first frame received and steps 228, 230 and 232 are performed in which an entry is made in the reassembly table 180. Specifically, in step 228, the Sequence Count register value is decremented by one and the decremented value is inserted into the sequence count field 184a of the first entry 182a in the reassembly table 180. In step 230, the PREVIOUS DATA pointer from the Previous Data Pointer register 136 is inserted into the data pointer field 186a of the memory entry 182a and in step 232, the value stored in the Data Length register 128 is inserted into the data length field 188a of the memory entry 182a.

Thereafter, the current context registers 118 are reset in preparation for the next received frame 50a-n in steps 234, 236, 238, and 240. Note that where the Frame counter 134 is determined to be equal to one in step 226, step 234 is performed directly after step 226. In step 234, the Sequence Count register 122 is set to a value of one plus the sequence count of the received frame, for comparison to the sequence count of a subsequently received frame 50a-n. In step 236, the PREVIOUS DATA pointer is made to point to the same location as the CURRENT DATA pointer. In step 238, the Data Length register 128 is set to by the length of data 56a associated with the received frame 50a. And, in step 240, the CURRENT DATA pointer is incremented by the length of the received frame data 56a, to point to the next available memory location for storing incoming frame data.

In step 242, it is determined whether all of the frames 50a-n of the packet 34 have been received. Specifically, the value in the Frame counter 134 is compared to one more than the difference between the highest sequence count and the lowest sequence count, as determined from bit nineteen of the F_CTL field 82 of the header of the received frame and the SOF field 52 of the received frame, respectively. A Frame counter value equal to one more than the highest

sequence count minus the lowest sequence count indicates that all of the frames 50a-n of the packet 34 have been received. The above-described process is repeated, starting at step 210, until the last frame associated with the packet 34 is received.

Once the Frame counter 134 indicates that all of the frames 50a-n for a particular packet 34 have been received, the last entry 182n is made in the reassembly table 180 in steps 244, 246, and 248. Specifically, in step 244, the Sequence Count register 122 is decremented by one and the decremented value is inserted into the sequence count field 184n of the last entry 182n. In step 246, the PREVIOUS DATA pointer from the Data Pointer register 136 is inserted into the data pointer field 186n of the last memory entry 182n and, in step 248, the value stored in the Data Length register 128 is inserted into the data length field 188n of the last memory entry 182n.

In order to illustrate the operation of the port 38 in reordering non-consecutive frames, consider two cases where a packet 34', including six frames 50a-50f having respective sequence counts of 0-5, consecutively, and transmitted in the order shown in FIG. 7A, is received by the port 38. Specifically, consider first the case where the frames 50a-50f are received by the port 38 in the following order: frame 50a, frame 50b, frame 50f, frame 50c, frame 50d and frame 50e, as shown in FIG. 7B. The frames 50a-50f have respective data fields 56a-56f with lengths as indicated in the data fields 56a-56f in FIGS. 7A-C. FIG. 8 shows a portion of the system memory 158 at which the data 56a-56f, associated with the frames 50a-50f, respectively, of FIG. 7B are stored. FIG. 9 shows the reassembly table 180' generated by the Receive Frame Manager 100 and stored in local memory 126 in response to receipt of frames 50a-50f in the order illustrated in FIG. 7B.

When the first frame 50a is received, the registers 118 of the current context 178 are in their initialized states (i.e., because the tag generated in response to receipt of frame 50a was found not to match a tag in the Receive cache 150 so the system processor 156 initialized the current context 178 in step 212). Thus, the Sequence Count register 122 is at a value of zero, the Data Length register 128 is at a value of zero and the Frame counter 134 is at a value of zero. Also, the PREVIOUS DATA pointer points to the location 250 (FIG. 8) to which the CURRENT DATA pointer is initialized by the system processor 156.

The data 56a associated with the received frame 50a is stored in system memory 158, starting at the location 250 pointed to by the CURRENT DATA pointer in step 214. In step 216, the Frame counter 134 is incremented by one, to a value of one. When the sequence count of the received frame 50a (i.e., zero) is compared to the sequence count of zero contained in the register 122 in step 218, it is determined that the two sequence counts are equal, thereby indicating that the frame 50a has been received in the same order in which it was transmitted.

In accordance with the process discussed above, there will not be an entry made in the reassembly table 180' in response to receipt of frame 50a since the frame 50a is received in the same order in which it appears in the transmitted frame 34, namely first. Stated differently, frame 50a is the first expected frame. The registers 118 of the context 178 are incremented in steps 220-224. Specifically, in step 220, the Sequence Count register 122 is incremented by one, to a value of one. In step 222, the Data Length register 128 is incremented by the length of the data 56a associated with the received frame 50a, to a value of 512 bytes. And, in step 224,

13

the CURRENT DATA pointer is incremented by the length of the data 56a associated with the received frame 50a (i.e., in the illustrative example, by 512 bytes, to memory location 252).

Thereafter, in step 242 (FIG. 6B), it is determined that the Frame counter 134, now at a value of one, is not equal to one more than the difference between the highest sequence count of five and the lowest sequence count of zero. Thus, since the received frame 50a is not the last frame of the packet 34', the process is repeated, starting at step 210.

When the second frame 50b associated with the packet 34' is received, generation of the tag associated with the incoming frame 50b and comparison of the generated tag to the current context 178 in the Receive Frame Manager 100 in step 212 reveals that the subsequently received frame 50b is associated with the same packet 34' as the previously received frame 50a. Thus, the contents of the context 178 are maintained for further processing in conjunction with frame 50b. The data 56b associated with the frame 50b is stored in the system memory 158 starting at the memory location 252 pointed to by the CURRENT DATA pointer in register 138 in step 214. In step 216, the Frame counter 134 is incremented by one, to indicate that two frames 50a and 50b of the packet 34' have been received.

The determination of process step 218 reveals that the sequence count of one, of the received frame 50b equals the current value of one in the Sequence Count register 122. Thus, no entry is made in the reassembly table 180' in response to receipt of the frame 50b since frame 50b is consecutive with respect to the previously received frame 50a. Thereafter, the registers 118 are incremented in steps 220-224. Specifically, in step 220, the Sequence Count register 122 is incremented by one to a value of two. In step 222, the Data Length register 128 is incremented by the length of the received frame data 56b. In the illustrative case, the Data Length register 128 is incremented by 256 bytes, to 768 bytes. In step 224, the CURRENT DATA pointer is also incremented by the length of the received frame data 56b, to point to memory location 254 in FIG. 8.

In step 242, it is determined that the received frame 50b is not the last frame of the packet 34' to be received, thereby causing the process to be repeated a third time, starting at step 210. Since the next frame received 50f is non-consecutive with respect to the previously received frame 50b, it is determined in step 218 that the sequence count of five, associated with the received frame 50f, is not equal to the sequence count value of two currently stored in the register 122. Thus, a first entry 182a' will be provided in the reassembly table 180' in response to non-consecutive frame 50f as follows.

Since the value in the Frame counter 134 is not equal to one (as determined in step 226), the Sequence Count register 122 is decremented by one, to a value of one and this decremented value of one is inserted into the sequence count field 184a' of the first entry 182a' of the table 180' in step 228, as shown in FIG. 9. The data pointer field 186a' of the entry 182a' contains the PREVIOUS DATA pointer from register 136 which still points to location 250 and the data length field 188a' of the entry 182a' contains the value from the Data Length register 128, of 768 bytes.

Once the first memory entry 182a' is completed, the registers 118 are readied for receipt of a subsequent frame, in steps 234-240, by setting the Sequence Count register 122 to one plus the sequence count of the received frame (i.e., to a value of six), by moving the CURRENT DATA pointer to the PREVIOUS DATA pointer location 254, by setting the

14

Data Length register 128 to the length of data 56f of associated with the third received frame 50f of 1024 bytes, and by incrementing the CURRENT DATA pointer maintained in register 138 by 1024 bytes (i.e., the length of the data 56f associated with the frame 50f) to point to location 256 at byte 1792.

Frame 50c is the fourth frame to be received. Since the Sequence Count register 122 is now at a value of six, it is determined that the sequence count of two of the received frame 50c is not equal to the value in the Sequence Count register 122. Thus, in response to receipt of frame 50c, there will be a second entry 182b' provided in the reassembly table 180'. In step 226, it is determined that the Frame counter 134, now at a value of four, is not equal to one. Thus, the Sequence Count register 122 is decremented by one, to a value of five, and this value of five is inserted into the sequence count field 184b' of the second entry 182b' in the reassembly table 180'. In step 230, the PREVIOUS DATA pointer, pointing to location 254, is inserted into the data pointer field 186b' of the second entry 182b' in the reassembly table 180'. The second table entry 182b' is completed in step 232 by the insertion of the value in the Data Length register 128, in this case, a value of 1024 bytes, as shown in FIG. 9.

Steps 234-240 are next performed in which the registers 118 are again reset in preparation for receipt of a subsequent frame. Specifically, the Sequence Count register 122 is set to a value of one more than the sequence count of the received frame, or in this case a value of three. In step 236, the PREVIOUS DATA pointer is modified to point to memory location 256 to which the CURRENT DATA pointer currently points. In step 238, the Data Length register 128 is set to the length of the data 56c associated with the received frame 50c. And, in step 240, the CURRENT DATA pointer is incremented by the length of the received data 50c, here 512 bytes, to point to memory location 258. Thereafter, in step 242, it is determined that the received frame 50c is not the last frame associated with the packet 34', causing the process to be repeated again, starting at step 210.

When the fifth frame 50d is received, the Frame counter 134 is incremented by one, to a value of five. It is determined, in step 218 that frame 50d is consecutive with respect to previously received frame 50c. Thus, no entry is made in table 180' in response to frame 50c and the registers 118 are reset in steps 220-224. Specifically, the Sequence Count register 122 is incremented by one in step 220 to a value of four and the Data Length register 128 is incremented by length of data 56d associated with the received frame 50d (i.e., by 2048 bytes) to a value of 2560 bytes in step 222. Also, the CURRENT DATA pointer is incremented by the length of the data 56d to a point to memory location 260 in step 224.

Finally, frame 50e is received and processed in the same manner as the preceding frame 50d, since frame 50e is consecutive with respect to the previously received frame 50d. Thus, since the sequence count of four of the received frame 50e is equal to the current value of four in the Sequence Count register 122, no entry is made in the reassembly table 180' in response to frame 50d. The Sequence Count register 122 is incremented by one, to a value of five, in step 220 and the Data Length register 128 is incremented by the length of data 56e associated with the received frame 50e in step 222, to a value of 2688 bytes. Also, in step 224, the CURRENT DATA pointer is incremented by 128 bytes to point to memory location 262.

Thereafter, in step 242, it is determined that the Frame counter 134 is at a value equal to one more than the highest

sequence count minus the lowest sequence count, or six, the total number of frames in the illustrative packet 34'. In response to this determination, Sequence Count register 122 is decremented by one, to a value of five, and this value is inserted into the sequence count field 184c' of the last memory entry 182c' in step 244. Additionally, the PREVIOUS DATA pointer location 256 is inserted into the last memory entry 182c' in step 246 and the Data Length register value of 2688 is inserted into memory entry 182c' in step 248.

With this arrangement, a reassembly table 180' has been generated including an entry for each non-consecutively received frame 50f and 50c, and for the last frame 50e of the packet 34' to be received. The table 180' is now ready for processing by the system processor 156 to reorder the frames 50a-50f so as to reassemble the packet 34'. Specifically, the system processor 156 processes each of the entries 182a'-182c' in ascending order of sequence count. Thus, entry 182a' is processed first, followed by entry 182c' and finally by entry 182b'.

Processing of an entry is achieved by accessing memory locations starting at the pointer contained in the pointer field for the entry and specifically, accessing a number of memory locations specified by the data length field of the entry. Thus, the packet 34' is reassembled by processing entry 182a' by reading 768 bytes of memory starting at location 250, followed by reading 2688 bytes of memory starting at location 256 and finally, by reading 1024 bytes of memory starting at location 254. In this way, the reassembled packet 34' is provided.

Referring to FIG. 7C, consider the case where the frames 50a-50f of the transmitted packet 34' (FIG. 7A) are received by the port 38 in the following order: frame 50b, frame 50a, frame 50f, frame 50c, frame 50d and frame 50e, as shown in FIG. 7C. This example illustrates the case where the first frame of a packet to be received is not the first frame of the packet as transmitted (i.e., where the received frame is not the expected first frame). FIG. 10 shows a portion of the system memory 158 at which the data 56a-56f, associated with frames 50a-50f, respectively, are stored and FIG. 11 shows the reassembly table 180" generated by the Receive Frame Manager 100 and stored in local memory 126 in response to receipt of frames 50a-50f in the order illustrated in FIG. 7C.

Prior to receipt of frame 50b, the registers 118 of the context 178 are initialized in steps 202-208, as described above. Thus, the Sequence Count register 122 is at a value of zero (assuming that the frame having a sequence count of zero is the first expected frame), the Data Length register 128 is at a value of zero, the PREVIOUS DATA pointer register 136 points to the same memory location to which the CURRENT DATA pointer is initialized. When frame 50b is received, the frame data 56b is stored at the initialized memory location 270 pointed to by the CURRENT DATA pointer and the Frame counter 134 is incremented by one, to a value of one. In step 218, it is determined that the sequence count of the received frame 50b (i.e., a sequence count of one) is not equal to the value of zero stored in the Sequence Count register 122. Thus, process step 226 is next performed in which it is determined that the Frame counter 134 is at a value of one. This determination corresponds to the special case where the first frame to be received is not the expected first frame (i.e., is not the first frame of the packet 34' as transmitted). In this case, no entry is made in the memory table 180" for this first, non-consecutively received frame 50b. Rather, after step 226, steps 243-240 are performed, in which the registers 118 are reset in preparation for receipt of a subsequent frame.

Specifically, in step 234, the Sequence Count register 122 is set to a value of one greater than the sequence count of the received frame, or in this case to a value of two. In step 236, the CURRENT DATA pointer is moved to the PREVIOUS DATA pointer location 270 and in step 238, the Data Length register 128 is set to the length of the received frame data 56b (i.e., to 256 bytes). In step 240, the CURRENT DATA pointer is incremented by the length of the received frame data 56c, to memory location 272. Thereafter, in step 242, it is determined that the Frame counter value of one is not equal to one more than the difference between the highest sequence count and the lowest sequence count, thereby indicating that the last frame of the packet has not been received. Thus, the process is repeated, starting at step 210.

Once the next frame 50a is received, it is determined that the tag associated with the received frame 50a matches the tag associated with previously processed frame 50b. Thus, in step 212, the contents of the current context 178 are maintained for processing frame 50a. In step 214, the data 56a, associated with frame 50a, is stored starting at the memory location 272 pointed to by the CURRENT DATA pointer. In step 216, the Frame counter 134 is incremented to a value of two. Thereafter, in step 218, it is determined that the sequence count of zero of the received frame 50a is not equal to the sequence count value of two stored in the Sequence Count register 122. Thus, an entry 182a" will be made in the memory table 180" in response to receipt of frame 50a.

In step 226, it is determined that the Frame counter 134 is not equal to one, causing steps 228-232 to be performed, in which the first entry 182a" is made in the memory table 180". Specifically, in step 228, the Sequence Count register 122 is decremented by one, to a value of one, and that value is inserted into the first entry 182a" in table 180". In step 230, the PREVIOUS DATA pointer pointing to memory location 270 is inserted into the entry 182a" and in step 232, the value of 256 bytes stored in the Data Length register 128 is inserted into the entry 182a".

In steps 234-240, again, the registers 118 are reset for receipt of the next frame. Step 234 causes the Sequence Count register 122 to be set to a value of one greater than the sequence count of the received frame, or to a value of one. The PREVIOUS DATA pointer is made to point to memory location 272 in step 236 and the Data Length register 128 is set to 512 bytes, the length of the data 56a associated with frame 50a. In step 240, the CURRENT DATA pointer is incremented by 512 bytes, to point to memory location 274. Since received frame 50a is not the sixth, and last frame of the packet to be received, the process is repeated.

Thereafter, frame 50f is received and the data 56f associated with the received frame is stored in memory starting at location 274, to which the CURRENT DATA pointer currently points. The Frame counter 134 is incremented to a value of three in step 216 and, in step 218, it is determined that the sequence count of five of the received frame 50f is not equal to the sequence count of one in the register 122. Thus, since frame 50f is received non-consecutively with respect to the previously received frame 50a, a second entry 182b" is made in the table 180" in accordance with steps 228-234.

In step 228, the Sequence Count register 122 is decremented by one, to a value of zero and this decremented value is inserted into the second entry 182b" as shown in FIG. 11. Also, the PREVIOUS DATA pointer location 272 is inserted into the second entry 182b" in step 230 and the value of 512 bytes in the Data Length register 128 is inserted into the

17

second entry 182b" in step 232. Since the subsequently received frames 50c, 50d, and 50e are received in the same order as described above in conjunction with the example of FIGS. 7B, 8, and 9, the process by which entries 182c" and 182d" are made into table 180" is identical to the process discussed above in which entries 182b' and 182c' are made into the table 180'.

The reassembly table 180" shown in FIG. 11 is processed in the same manner as table 180' discussed above. Specifically, the system processor 156 processes entries 182a"-182d" starting with the entry having the lowest sequence count. Thus, entry 182b" is processed first followed by entry 182a", entry 182d" and finally 182c". The packet 34' of FIG. 7A is reassembled by reading 512 bytes from memory starting at location 272, followed by 256 bytes starting at memory location 270, 2688 bytes starting at memory location 276 and finally, 1024 bytes starting at memory location 274.

It will be appreciated that, in the case where a packet 34 is received in the same order as it is transmitted, the reassembly table 180 for that packet 34 will contain only a single entry, corresponding to the last frame of the packet 34 to be received. Such an entry includes the sequence count of the last received frame, a pointer to the location in memory 158 beginning at which the first frame of the packet 34 is stored, and a data length value indicating the cumulative length of data of each of the frames 50a-n of the packet 34.

It will be appreciated that various options exist for implementing the port 38, such as by integrating the port components on an ASIC. It will also be appreciated that the order of the process steps of FIGS. 6A and 6B may be varied somewhat without departing from the spirit of the invention.

It will also be appreciated that, although the reassembly tables are described hereinabove as being stored in local memory 126 various memory configurations can be used without departing from the spirit of the invention. For example, the reassembly tables may be stored in the system memory 158. Furthermore, the entries in a reassembly table need not be stored in a contiguous block of memory locations.

Referring to FIG. 12, an alternate node embodiment 300 is connected to the switch fabric 10. Recall from FIG. 1 that the switch fabric 10 includes a plurality of switches 12-20 interconnecting a plurality of nodes. In the network of FIG. 12, fabric 10 interconnects nodes 22 and 300. Node 22 includes a port 36 of the type discussed above in conjunction with FIG. 3. Since node 22 is associated with a single switch fabric channel 301, port 36 is used interchangeably with node 22. Node 300 on the other hand is referred to as a shared node 300 and includes multiple ports 304₁-304_N adapted for connection to multiple respective channels 302₁-302_N of the switch fabric 10, as shown. The number of channels 302₁-302_N serviced by the node 300 and thus, the number of ports 304₁-304_N of the node, may be readily varied and is selected generally, in accordance with factors such as the node memory capacity, available circuit board area for the node, and the available pin count where the node 300 is incorporated onto an integrated circuit, such as an ASIC, as will become apparent. Preferably, the number of channels 302₁-302_N serviced by the shared node 300 and thus, the number of ports 304₁-304_N, is a binary multiple (i.e., 2, 4, 8, 16, ...).

While the ports 36, 38, 40 of FIG. 1 are described as including the circuitry of FIG. 3, in the shared node 300, the term "node" refers to the entire switch fabric interconnection circuit and the ports 304₁-304_N refer to the terminal, or

18

physical connection between the respective channel 302₁-302_N and the node 300. Thus, each port 304₁-304_N is associated with a respective channel 302₁-302_N. Further, each port 304₁-304_N has a respective node client system 306, 308 and 310 associated therewith and connected to the shared node 300 via a system bus 316. Node clients 306, 308, 310 may be of various types, such as the illustrated host computer 306, data storage device 308, or network interface 310. Preferably, the shared node 300 is integrated onto an ASIC using 0.35 micron technology.

Referring also to FIG. 13, the shared node 300 includes a portion 314₁-314_N dedicated to each port 304₁-304_N and a common portion shared by each of the ports 304₁-304_N. The common portion of the node 300 includes a frame handler 320, a frame prioritization circuit 330 and a frame routing circuit 340. The shared node 300 is adapted for dynamically receiving frames of a packet from one or more source nodes over one or more respective channels 302₁-302_N and for transmitting frames of a packet to one or more destination nodes over one or more channels 302₁-302_N. The shared node 300 operates dynamically, in the sense that frames can be received by one or more ports simultaneously with respect to one another, while another frame is being processed by the frame handler 320. Thus, when the frame handler 320 processes a frame received at a particular port, the other ports are not rendered "busy" but rather, are capable of receiving frames.

The portion of node 300 dedicated to each port 304₁-304_N includes a receiver and transmitter circuit 314₁-314_N associated with a respective port 304₁-304_N of the node 300. Each of the receiver and transmitter circuits 314₁-314_N includes circuitry for receiving frames at the respective port 304₁-304_N and for transmitting frames to the respective port 304₁-304_N, as will be discussed further in conjunction with FIGS. 14 and 14A.

The frame handler 320 processes frames received by each of the ports 304₁-304_N. More particularly, the frame handler 320 performs various types of frame control and processing functions, such as reordering frames of a packet received out-of-order, in the manner described above. The frame handler 320 includes a memory manager and memory 324 and a system interface 328. The memory manager and memory 324 provide control and temporary storage of incoming frame data and the system interface 328 permits communication between the node 300 and the node clients 306, 308, 310 via bus 316.

The receiver portion of each receiver and transmitter circuit 314₁-314_N separates the header 54a-n (FIG. 2) of received frames 50a-n from the frame data 56a-n and temporarily stores the headers 54a-n and data 56a-n thus separated. The headers 54a-n of the received frames are selectively forwarded by the frame prioritization circuit 330 to the frame handler 320 via signal lines 332₁-332_N and 333 and the data 56a-n of the received frames are transferred to the memory manager and memory 324 via signal lines 334₁-334_N.

The frame prioritization circuit 330 selects a frame header temporarily stored in one of the receiver and transmitter circuits 314₁-314_N in accordance with one of a plurality of port prioritization schemes and may alternatively be referred to as a multiplexer 320. In one embodiment, each of the ports 304₁-304_N is assigned a priority at system initialization based on a predetermined criteria. For example, the predetermined criteria may be based on the type of node client associated with the respective port and/or the anticipated frame traffic on the respective port. In operation, the frame

prioritization circuit 330 initially forwards one or more headers from the receiver and transmitter circuit 314₁-314_N associated with the port 304₁-304_N assigned the highest priority. Thereafter, the selection of subsequent frame headers for forwarding to the frame handler 320 is based on a descending order of port priority.

In accordance with another prioritization scheme, received frame headers 54a-n are selected for forwarding to the frame handler 320 in accordance with a predetermined order of ports, such as in a "round robin" fashion, in which the ports 304₁-304_N are polled in a consecutive repetition, from port 304₁ to port 304_N. Alternatively, the ports 304₁-304_N may be polled on a first frame received, first frame forwarded basis.

A further, adaptive prioritization scheme is based on which port 304₁-304_N has received the most frames, the headers of which are temporarily stored by the respective receiver and transmitter circuit 314₁-314_N. More particularly, the frame prioritization circuit 330 accesses the receiver and transmitter circuit 314₁-314_N with the most stored frame headers for forwarding one or more such frame headers to the frame handler 320. Subsequent receiver and transmitter circuits 314₁-314_N are then accessed in a descending order of the most temporarily stored frame headers.

Once a frame header 54a-n is forwarded to the frame handler 320, the respective frame 50a-n is processed by the frame handler 320. In response to a request by a node client to transmit a frame, the frame is processed by the frame handler 320 and the processed frame is routed to the frame routing circuit 340 which receives the processed frame via signal line 338, determines the destination node identified by the frame and forwards the frame a port 304₁-304_N associated with the destination node. The frame routing circuit 340 may alternatively be referred to as a demultiplexer and will be described further in conjunction with FIG. 16.

Referring also to FIGS. 14 and 14A, the receiver and transmitter circuit 314₁-314_N associated with each port 304₁-304_N includes a deserializer 341₁-341_N, a serializer 342₁-342_N, a decoder 360₁-360_N, an encoder 370₁-370_N, a PSR (Primitive Sequence Recognition) circuit 362₁-362_N, a PSG (Primitive Sequence Generation) circuit 372₁-372_N, an arbitrated loop (AL) circuit 344₁-344_N, a Frame Receive circuit 364₁-364_N, a Frame Transmit circuit 376₁-376_N, a Frame Data buffer, or FIFO 350₁-350_N, a Frame Header buffer, or FIFO 352₁-352_N, and a Frame Transmit buffer, or FIFO 354₁-354_N, as shown. Since each receiver and transmitter circuit 314₁-314_N is identical, illustrative circuit 314₁ associated with port 304₁ will be described in greater detail.

Each channel 302₁-302_N (FIG. 12) has a pair of serial signal lines associated therewith, a first one of which carries frames serially to the respective port 304₁-304_N from the switch fabric and a second one of which carries frames serially from the port 304₁-304_N to the switch fabric 10. The deserializer 341₁ converts serial frame data received at port 304₁ into parallel frame data. Similarly, the serializer 342₁ converts the parallel frame data into a single serial signal.

Although preferably, the deserializer 341₁ and serializer 342₁ are incorporated onto the shared node ASIC 300, the deserializer and serializer may alternatively be external to the node. In this case however, the pin count of the ASIC is increased in accordance with the number of parallel signal lines connected between the external deserializer and the ASIC and between the external serializer and the ASIC.

The receiver portion of the receiver and transmitter circuits 314₁-314_N includes, in addition to the deserializer

341₁-341_N, the circuitry contained in the Frame receiver 110 of FIG. 3; namely, a decoder 360₁-360_N, a PSR circuit 362₁-362_N, a Frame Receive circuit 364₁-364_N, a Frame Data FIFO 350₁-350_N and a Frame Header FIFO 352₁-352_N. Considering illustrative circuit 314₁, the decoder 360₁ provides the necessary 8B/10B decoder functions by which received transmission characters (i.e., each character being ten bits) are decoded into eight-bit codes. The PSR circuit 362₁ decodes primitive sequences, which are three consecutive ordered Fibre Channel sets of the same type. An ordered set is a four character combination of data and special transmission characters. Illustrative types of ordered sets include frame delimiters which immediately precede or follow the contents of a frame and primitive signals which are ordered sets designated to have special meaning, such as an idle.

A Fibre Channel Arbitrated Loop specification (FC-AL) Rev. 4.34 X3T11 Project 960D dated Sep. 18, 1994, which is incorporated herein by reference, describes both switch fabric and arbitrated loop configurations. In an arbitrated loop configuration, nodes are daisy-chain connected to form a loop, with the transmitter circuitry of one node connected to the receiver circuitry of another. Nodes connected in an arbitrated loop configuration arbitrate for access to the loop by passing arbitration primitives around loop. Once a node obtains access to the loop, thereby becoming a source node, a destination node is "opened" and the transmission between the source node and destination node becomes a point to point transmission, with intermediate nodes positioned between the source and destination nodes on the loop passing data through to the next port on the loop.

The arbitrated loop circuit 344₁-344_N associated with each port 304₁-304_N is operable to pass frames through the respective port 304₁-304_N when an arbitrated loop topology is utilized and the respective port 304₁-304_N is an intermediate port on the loop. For example, when port 304₁ is an intermediate port in an arbitrated loop configuration, the arbitrated loop circuit 344₁ passes received frames through, for transmission to the next port on the loop.

The Frame Receive circuit 364₁ encapsulates the necessary logic to properly receive the Fibre Channel frames, verifies the transmitted CRC, and routes the frame data to a Local Memory Manager 382 of the frame handler 320. More particularly, the Frame Receive circuit 364₁ separates the frame header 54 from the data 56 (FIG. 2) for further processing. The frame data 56 is transferred to the Frame Data FIFO 350₁ for further transfer to a local memory 380 via the Local Memory Manager 382. More particularly, the Frame Receive circuit 364₁ passes control/delimiter data to the Local Memory Manager 382, which in turn transfers the frame data from Frame Data FIFO 350₁ to local memory 380. The frame header 54 is transferred to the Frame Header FIFO 352₁ for temporary storage.

The Frame Header FIFOs 352₁-352_N are connected to the frame prioritization circuit 330 via signal lines 332₁-332_N and 356₁-356_N, as shown. Frame headers are forwarded to the frame prioritization circuit 330 via signal lines 332₁-332_N and control signals described below are transferred via signal lines 356₁-356_N. The frame prioritization circuit 330 forwards selected received frame headers to the frame handler 320 via signal line 333, as described below in conjunction with FIG. 15.

The frame handler 320 includes much of the same circuitry as the illustrative, single channel port 38 shown in FIG. 3, such as a local memory 380, a Local Memory Manager 382, a Frame Preprocessor 384, a Receive Frame

manager 386, a Receive cache 388, a DMA 390, a Transmit Frame manager 392, a Transmit cache 394 and the System Interface 328. Also included in the frame handler 320 are a Write Data FIFO 396 and a Read Data FIFO 398 associated with the Receive cache 388, a Write Data FIFO 400 and a Read Data FIFO 402 associated with the DMA 390, and a local memory Write Data FIFO 406 and Read Data FIFO 408, as shown. The circuitry of the frame handler 320 operates in substantially the same manner discussed above in conjunction with FIG. 3. However, whereas the Frame preprocessor 120 of FIG. 3 is responsive to frame headers from the single Frame Header FIFO 114, the Frame preprocessor 384 is responsive to frame headers from the plurality of Frame Header FIFOs 352₁-352_N forwarded to the Frame preprocessor 384 by the frame prioritization circuit 330.

More particularly, a frame header forwarded to the frame handler 320 by the frame prioritization circuit 330, is specifically forwarded to the Frame preprocessor 384. The Frame preprocessor 384 checks the CRC field 58 (FIG. 2A) of the received frame header to determine whether or not the frame will be discarded if the frame was not accurately received, maintains error counters, and determines whether the destination node specified in the header field 76 of the received frame matches that of the receiving port, as described above. After processing a received frame header, the Frame preprocessor 384 forwards the frame header to the Receive Frame Manager 386.

The Receive Frame Manager 386 is described above and shown in FIG. 4 and controls the transfer of received frame data 56 from the local memory 380 to the system memory associated with each node client system by programming the DMA 390 to move data 56 out of the local memory 380 and transfer such data to the system memory in response to receipt of the respective frame header 54 by the Receive Frame Manager 386. The DMA 390 achieves this transfer of data 56 via Write Data FIFO 400, Read Data FIFO 402, the Local Memory Manager 382, and the System Interface 328.

The Transmit Frame Manager 392 programs the DMA 390 to read data 56 out of the system memory and forward the data, along with a corresponding header, to the frame routing circuit 340 in response to a request from a node client. The header for transmission to the destination node may be generated by the Transmit Frame Manager 392 or, alternatively, may be provided directly by the host computer 306 or other node client.

The Receive cache 388 is provided in communication with the Receive Frame Manager 386 and the Local Memory Manager 382 via the Write Data FIFO 396 and the Read Data FIFO 398 and stores a context associated with each packet processed by the node 300, as described above in conjunction with FIGS. 4 and 5. When a frame header is received by the Receive Frame Manager 386, it is determined whether the class of transmission service requires that an acknowledgement signal be returned to the source node. In the event that an acknowledgement signal is required, the control state machine of the Receive Frame Manager 386 (labelled 174 in FIG. 4) generates the acknowledgement signal and forwards such acknowledgement signal to the Transmit Frame Manager 392. The receive sequence state machine of the Receive Frame Manager 384 (labelled 170 in FIG. 4) generates a tag associated with a received frame header for comparison to contexts stored in the Receive cache 388. If the generated tag matches one stored in the Receive cache 388, then the contents of the matching context are loaded from the cache 388 into the current context of the Receive Frame Manager (labelled 178 in FIG. 4). Alternatively, if the generated tag does not match a

context stored in the Receive cache 388, then the receive sequence state machine 170 receives initialized contents for the current context 178 from the node client associated with the port on which the processed frame was received. Similarly, the Transmit cache 394 is provided in communication with the Transmit Frame Manager 392 and the Local Memory Manager 382 via the Write Data FIFO 406 and the Read Data FIFO 408, as shown. The Transmit cache 394 functions in a manner similar to the Receive cache 388, albeit in the transmission of frames as opposed to the receipt of frames.

In response to a request by a node client 306, 308, 310, frames are processed by the frame handler 320 and are forwarded to the frame routing circuit 340 for transmission to other nodes of the fabric 10. More particularly, the processed frame is provided to the frame routing circuit 340 via signal line 338. The frame routing circuit 340 routes processed frames to a port 304₁-304_N for transmission to one or more destination nodes via the switch fabric 10, in a manner described further in conjunction with FIG. 16.

The transmitter portion of each receiver and transmitter circuit 314₁-314_N (i.e., the circuitry included in the Frame Transmitter 146 of FIG. 3) includes an encoder 370₁-370_N, a PSG circuit 372₁-372_N, a Frame Transmit circuit 376₁-376_N, and a Frame Transmit FIFO 354₁-354_N. Considering illustrative port 304₁, the Frame Transmit circuit 376₁ contains logic for assembling Fibre Channel frames based on the frames stored in the Frame Transmit FIFO 354₁, for calculating the CRC and for signalling the PSG circuit 372₁ to transmit idles between frames of a packet (see FIG. 2A). The PSG circuit 372₁ generates the appropriate primitive sequences for transmission on the respective channel 304₁ and inserts idles under the command of the Frame Transmit circuit 376₁. The encoder 370₁ generates ten-bit transmission characters for transmission over the switch fabric 10.

Referring also to FIG. 15, a block diagram of the frame prioritization circuit 330 is shown. The frame prioritization circuit 330 includes a switch 430 having a plurality of input terminals 432₁-432_N and an output terminal 434. The switch 430 is adapted for being in one of "N" positions in which the output terminal 434 is connected to a respective one of the "N" input terminals 432₁-432_N. Also provided in the frame prioritization circuit 330 is a control circuit 438 for controlling the position of the switch 430.

The control circuit 438 implements one or more of the port prioritization schemes described above for selecting a frame header received by one of the ports 304₁-304_N to be forwarded to the frame handler 320 for processing of the respective frame. In one such prioritization scheme, each port 304₁-304_N has a predetermined priority assigned thereto at system initialization. The priority assigned to each port 304₁-304_N may be based on various criteria, such as the type of node client associated with the port or the signal processing rate thereof or the anticipated frame traffic on the port. For example, it may be advantageous to assign a higher priority to the port 304₁ associated with the host computer 306 than to the port 304₂ associated with the data storage device 308.

A port identifier and the respective port priority are stored in a memory portion of the control circuit 438 for use in selecting a received frame header for forwarding to the frame handler 320. With this arrangement, the control circuit 438 polls the Frame Header FIFO 352₁-352_N associated with each port 304₁-304_N in the order of priority assigned to the respective port 304₁-304_N. For example, consider the case where node 300 includes ports 304₁, 304₂, and 304_N where

N=3, as shown in FIG. 12, and port 304₁ is assigned the highest priority, port 304₃ is assigned the second highest priority and port 304₂ is assigned the lowest priority. The control circuit 438 initially positions switch 430 to connect input terminal 432₁ to output terminal 434. With switch 430 in this position, one or more frame headers stored in Frame Header FIFO 352₁ are forwarded to the frame handler 320. Thereafter, control circuit 438 toggles switch 430 to a second position, in which input terminal 432₃ is connected to the output terminal 434 for forwarding one or more frame headers in FIFO 352₃ to frame handler 320. Finally, control circuit 438 causes switch 430 to move to a third position in which input terminal 432₂ is connected to output terminal 434 in order to forward one or more frame headers stored in Frame Header FIFO 352₂ to frame handler 320. This process of polling the Frame Header FIFOs 352₁-352_N associated with each port 304₁-304_N in ascending order of the priority assigned to the respective port continues until the port having the lowest priority is polled, following which the process is repeated. Note that all of the frame headers stored in a polled Frame Header FIFO 352₁-352_N may be forwarded to the frame handler 320 sequentially, so as to empty a selected FIFO each time the FIFO is connected to the frame handler 320 via switch 430 or, alternatively, only one, or another selected number, of frame headers thus stored may be forwarded each time a FIFO 352₁-352_N is connected to the frame handler 320 via switch 430.

In accordance with another prioritization scheme, the ports 304₁-304_N are accessed for forwarding stored frame headers in accordance with a predetermined order of ports 304₁-304_N. One such predetermined order is a "round robin" scheme, in which the ports are polled in accordance with the sequence of the ports 304₁-304_N, such that one or more frame headers received by port 304₁ are transferred to the frame handler 320, followed by one or more frame headers received by port 304₂, etc.

Another predetermined port order is based on the order in which frames are received at each port 304₁-304_N. In accordance with this prioritization scheme, the Frame Header FIFOs 352₁-352_N send empty/not empty signals on signal line 356₁-356_N to the frame prioritization circuit 330 indicative of whether a frame header has been received by the respective FIFO 352₁-352_N. This prioritization scheme operates on an interrupt basis in that, once an empty/not empty signal is received by the frame prioritization circuit 330 indicating that a frame header has been received, the frame prioritization circuit accesses that FIFO and transfers the received frame, or frames to the frame handler 320. In the case where more than one of the Frame Header FIFOs 352₁-352_N receives frames substantially simultaneously, causing the frame prioritization circuit 330 to receive more than one not empty signal substantially simultaneously, then the FIFOs thus receiving frames are polled in a round robin fashion.

In accordance with an adaptive port prioritization scheme, the control circuit 438 positions switch 430 in accordance with the number of frame headers currently stored in the Frame Header FIFO 352₁-352_N associated with each port 304₁-304_N. More particularly, in one embodiment, each Frame Header FIFO 352₁-352_N sends a level indication signal on respective signal line 356₁-356_N to the control circuit 438 indicative of how much data is currently stored in the FIFO 352₁-352_N. The control circuit 438 determines which FIFO 352₁-352_N has the most headers stored therein in response to the level indication signals and polls the FIFOs 352₁-352_N in descending order of the number of headers stored.

At system initialization, each port 304₁-304_N communicates to every other port connected to the fabric 10 information regarding the capacity of the respective Frame Header FIFO 352₁-352_N and Frame Data FIFO 350₁-350_N. In a Fibre Channel network, this communication is referred to as "login" and includes parameters such as how many frames a particular port can accept (i.e., the "credit"), the length of frame data that the port can accept, and the class of Fibre Channel service supported by the port, as described in the above-referenced Fibre Channel (FC-PH) specification. In operation, a transmitting, or source node uses the information about other nodes received during the initialization procedure to avoid overrunning the Frame Header FIFO and Frame Data FIFO of a destination node.

With the use of the shared node 300, the credit specified by a port at system initialization can be manipulated in accordance with the frame prioritization scheme employed by the frame prioritization circuit 330. More particularly, even though the Frame Header FIFO 352₁-352_N associated with each port 304₁-304_N has a fixed size, the credit can indicate that greater than actual buffer capacity is available if the frame prioritization scheme is tailored to ensure that the frames received at that particular port are rapidly forwarded to the frame handler 320. With this arrangement, a port can accommodate a higher frame traffic level than otherwise permissible by tailoring the frame prioritization scheme to provide the port with a higher rate of forwarding frames received by that port to the frame handler 320.

Referring also to FIG. 16, the frame routing circuit 340 includes a control circuit 450 which receives processed frames from the DMA 390 and Transmit Frame Manager 392 of the frame handler 320 via signal line 338 (FIGS. 14 and 14A). The control circuit 450 may include a buffer for storing frames received on line 338. The control circuit 450 includes a look-up table 454 which contains a list of destination node identifiers (i.e., the D_ID field 76 of the illustrative frame header 54 of FIG. 2B) for each node known to be connected to the switch fabric 10 and an identifier of a corresponding port 304₁-304_N of the shared node 300 on which a frame can be transmitted to each listed destination node. A switch 458 is provided with an input terminal 460 and a plurality of output terminals 462₁-462_N, each one corresponding to one of the ports 304₁-304_N of the shared node 300. Switch 458 is adapted for being in one of "N" positions, in which the input terminal 460 is connected to a respective one of the output terminals 462₁-462_N to route a processed frame from the frame handler 320 to one or more ports 304₁-304_N for transmission to one or more destination nodes.

In operation, when the control circuit 450 receives a processed frame on line 336 (i.e., either a data frame or an acknowledgement signal), the destination node identifier D_ID specified in the processed frame is located in the table 454. An identifier of the port 304₁-304_N associated with the located D_ID table entry is read from the table 454 and the control circuit 450 positions switch 458 accordingly. That is, circuit 450 positions switch 458 to connect the input terminal 460 to the output terminal 462₁-462_N associated with the port 304₁-304_N on which the processed frame is to be transmitted. In this way, the frame routing circuit 340 operates as a demultiplexer to connect the processed frame or acknowledgement signal at the input terminal 460 to at least one of the output terminals 462₁-462_N. Note that in broadcasted transmissions (i.e., transmission to more than one destination node), a processed frame may be transferred to multiple switch output terminals 462₁-462_N for transmission to multiple destination nodes via respective ports 304₁-304_N.

The shared node configuration described above in conjunction with FIGS. 12-16 is advantageous for several reasons. First, by sharing resources among multiple network channels 302₁-302_N, the cost of providing a node is reduced, as compared to the use of a dedicated port/node for each network channel, particularly when advanced processing technology, such as the preferred 0.35 micron technology, is used to fabricate the node 300. Additionally, the required circuit board surface area occupied by each node is reduced by use of the shared node 300, since disadvantageous circuit redundancy otherwise required is avoided. Furthermore, full utilization of resources shared among multiple ports, such as local memory, further reduces the effective cost of such shared resources to each port.

Additionally, the shared node 300 permits simultaneous processing of a frame and receipt of frames at the ports 304₁-304_N, thereby providing dynamic resource sharing. That is, as opposed to rendering one port busy when a frame received by another port is being processed, the shared node 300 permits dynamic receipt of frames by multiple ports 304₁-304_N when a frame is being processed by the frame handler 320.

Significantly, the shared node configuration permits enhanced node performance by permitting port frame traffic capability to be manipulated. That is, by tailoring the port prioritization scheme employed by frame prioritization circuit 330 to favor a particular port 304₁-304_N, that port can be permitted to accommodate a higher frame traffic level than otherwise permissible.

While the node in the illustrative embodiments is fabricated in the form of an ASIC, it will be appreciated that the circuitry may be fabricated on multiple ASICs or other utilizing other suitable circuit devices to implement the presently disclosed subject matter.

Having described the preferred embodiments, those skilled in the art will realize that many variations are possible which will still be within the scope and spirit of the claimed invention. Therefore, it is the intention to limit the invention only as indicated by the scope of the claims.

We claim:

1. A communications network for transmitting frames between nodes connected to the network, said network comprising:

(a) at least one node comprising:

- (i) a plurality of ports capable of receiving frames from another node connected to said network substantially simultaneously and transmitting frames to another node connected to said network;
- (ii) a frame handler for processing said frames received by said plurality of ports;
- (iii) a multiplexer for forwarding a first frame received by a selected one of said ports to said frame handler; and
- (iv) a demultiplexer for routing a second frame from said frame handler to at least a selected one of said plurality of ports; and

(b) a plurality of node clients in communication with said least one node, wherein each one of said node clients is associated with a corresponding one of said plurality of ports.

2. The communications network recited in claim 1 wherein said frame handler comprises a memory and a system interface for permitting communication between said frame handler and said plurality of node clients via a system bus.

3. The communications network recited in claim 1 wherein said node further comprises a plurality of receiver

and transmitter circuits, each one connected to a respective one of said plurality of ports.

4. The communications network recited in claim 1 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with a priority assigned to each of said plurality of ports, wherein a first one of said plurality of ports is assigned a highest priority and a second one of said plurality of ports is assigned a lowest priority.

5. The communications network recited in claim 1 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with a predetermined order of said plurality of ports.

6. The communications network recited in claim 1 wherein each of said frames received by said plurality of ports comprises a header and data and wherein said at least one node further comprises a plurality of frame header buffers, each one corresponding to a respective one of said plurality of ports, for storing the header of frames received by said respective ports.

7. The communications network recited in claim 6 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with the number of frame headers stored in each of said plurality of frame header buffers.

8. The communications network recited in claim 1 wherein said demultiplexer comprises a control circuit responsive to said second frame and a look-up table for determining a destination node specified in the header of said second frame and for routing said second frame to at least one of said plurality of ports associated with said destination node.

9. The communications network recited in claim 1 wherein each of said plurality of node clients is one of a host computer, a data storage device, and a network interface.

10. The communications network recited in claim 4 wherein said first one of said plurality of ports assigned the highest priority provides a credit to other nodes connected to said communications network indicating a number of frames that said first one of said plurality of ports is capable of receiving, wherein said indicated number of frames is greater than a number of frames that said respective frame header buffer is capable of storing.

11. A node for use in a communications network having a plurality of channels, said network permitting transmission of frames including a header between said node and another node connected to said network, said node comprising:

- a plurality of ports, each one associated with a respective one of said plurality of channels of said communications network;
- a plurality of frame header buffers, each one associated with a respective one of said plurality of ports, for storing the header of frames received by said respective port;
- a multiplexer for selecting a first frame stored in a one of said plurality of frame header buffers to be processed;
- a frame handler for processing said first frame; and
- a demultiplexer for routing a second frame designated for transmission to at least a selected one of said plurality of ports for transmission to said other node connected to said network.

12. The node recited in claim 11 further comprising a plurality of frame transmit buffers, each one associated with a respective one of said plurality of ports, for temporarily storing said second frame prior to transmission by said respective port, wherein said demultiplexer routes said sec-

27

ond frame to the respective frame transmit buffer associated with said selected port.

13. The node recited in claim 11 wherein said frame handler comprises a memory and a system interface for permitting communication between said frame handler and a plurality of node clients.

14. The node recited in claim 11 further comprising a plurality of receiver and transmitter circuits, each one associated with a respective one of said ports and being connected between said respective port and the respective frame header buffer.

15. The node recited in claim 11 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with a predetermined priority assigned to said plurality of ports.

16. The node recited in claim 11 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with a predetermined order of said plurality of ports.

17. The node recited in claim 11 wherein said multiplexer comprises a control circuit for selecting said first frame in accordance with the number of frame headers stored in each of said plurality of frame header buffers.

18. The node recited in claim 11 wherein said demultiplexer comprises a control circuit responsive to said second frame and a look-up table, for determining at least one of said plurality of ports associated with a destination node specified in said second frame and for routing said second frame to said at least one port.

19. A method for processing frames received by a node of a communications network, comprising the steps of:

receiving frames at a plurality of ports of said node, said ports being capable of receiving said frames substantially simultaneously with respect to one another;

28

selecting a first frame received by one of said plurality of ports for forwarding said first frame to a frame handler of said node for processing; and

routing a second frame designated for transmission to a destination node by said node to at least one of said plurality of ports associated with said destination node.

20. The method recited in claim 19 wherein said frame selecting step comprises the steps of:

assigning a priority to each of said plurality of ports based on a predetermined criteria; and

selecting said first frame in accordance with a descending order of port priority.

21. The method recited in claim 19 wherein said frame selecting step comprises the step of selecting said first frame in accordance with a predetermined order of said plurality of ports.

22. The method recited in claim 19 wherein said frame selecting step comprises the steps of:

determining the number of frames received at each of said plurality of ports; and

selecting said first frame in accordance with a descending order of the number of frames received by each of said plurality of ports.

23. The method recited in claim 19 wherein said frame routing step comprises the steps of:

determining a destination node identity specified by a header portion of said second frame;

accessing a look-up table to determine at least one of said plurality of ports associated with said destination node; and

routing said second frame to said at least one port.

* * * * *

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.